

• User Manual

CONTENTS

CHAPTER 1: INTRODUCTION 1

ABOUT THIS MANUAL	2
VES R2600 OVERVIEW	3
PERFORMANCE	3
GREENRAID	3
SERVICE AND SUPPORT	3
NEW FEATURES IN THE VES R2600	3
ARCHITECTURAL DESCRIPTION	4
VES R2000 SERIES MODEL LINE-UP	4
VES J2000 SERIES MODEL LINE-UP	4
FEATURES	6
CONTROLLER MODULE FEATURES	6
OPERATIONAL FEATURES	6
SYSTEM MANAGEMENT	7
GENERAL SPECIFICATIONS	9

CHAPTER 2: HARDWARE INSTALLATION 12

UNPACKING THE VES R2600	13
MOUNTING VES R2600 IN A RACK	14
INSTALLING PHYSICAL DRIVES	16
NUMBER OF DRIVES REQUIRED	16
DRIVE SLOT NUMBERING	16
INSTALLING YOUR DRIVES	17
MAKING MANAGEMENT AND DATA CONNECTIONS	18
FIBRE CHANNEL SAN	18
FIBRE CHANNEL DAS	20
FIBRE CHANNEL WITH JBOD EXPANSION	21
FIBRE CHANNEL SAN — NO SINGLE POINT OF FAILURE	22
iSCSI STORAGE AREA NETWORK (SAN)	24
iSCSI DIRECT ATTACHED STORAGE (DAS)	27
iSCSI WITH JBOD EXPANSION	28
MAKING SERIAL CABLE CONNECTIONS	29
CONNECTING THE POWER	30
LED BEHAVIOR	30
FRONT PANEL LEDs	31
DRIVE STATUS INDICATORS	32
CONTROLLER LEDs	32

CHAPTER 3: SETUP 35

SETTING-UP THE SERIAL CONNECTION	35
ABOUT IP ADDRESSES	36
VES R2600 DEFAULT IP ADDRESSES	36

CHOOSING DHCP OR A STATIC IP ADDRESS	36
ACCESSING THE MAC ADDRESS IN THE CLI	36
ACCESSING THE MAC ADDRESS IN THE CLU	37
SETTING-UP VES R2600 WITH THE CLI	37
MAKING SUBSYSTEM DATE AND TIME SETTINGS	37
VIRTUAL MANAGEMENT PORT SETTINGS	38
MAINTENANCE MODE SETTINGS	39

SETTING-UP VES R2600 WITH THE CLU	41
ACCESSING THE CLU QUICK SETUP MENU	42
MAKING SUBSYSTEM DATE AND TIME SETTINGS	42
VIRTUAL MANAGEMENT PORT SETTINGS	42
MAINTENANCE MODE SETTINGS	43
EXITING THE CLU	44

SETTING UP VES R2600 WITH WEBPAM PROE	44
LOGGING INTO WEBPAM PROE	44

CREATING DISK ARRAYS AND LOGICAL DRIVES	45
WIZARD - AUTOMATIC CONFIGURATION	46
WIZARD - ADVANCED CONFIGURATION	47
ENABLING LUN MAPPING AND MASKING	49
LOGGING OUT OF WEBPAM PROE	49
USING WEBPAM PROE OVER THE INTERNET	49

CHAPTER 4: MANAGE WITH WEBPAM PROE 50

LOGGING INTO WEBPAM PROE	50
CHOOSING THE DISPLAY LANGUAGE	51
PERUSING THE INTERFACE	51
LOGGING OUT OF WEBPAM PROE	52

VIEWING THE STORAGE NETWORK	53
LOGGING ONTO A SUBSYSTEM	53
FILTERING THE SUBSYSTEM LIST	53
REFRESHING THE LIST	53

MANAGING SUBSYSTEMS	53
VIEWING SUBSYSTEM INFORMATION	54
MAKING SUBSYSTEM SETTINGS	54
LOCKING OR UNLOCKING THE SUBSYSTEM	54
RESTORING FACTORY DEFAULT SETTINGS	55
CLEARING STATISTICS	55
SAVING A SERVICE REPORT	56
IMPORTING A CONFIGURATION SCRIPT	57
EXPORTING A CONFIGURATION SCRIPT	57
RESTARTING THE SUBSYSTEM	57
SHUTTING DOWN THE SUBSYSTEM	58
RESTARTING THE SUBSYSTEM AFTER A SHUTDOWN	58

MANAGING RAID CONTROLLERS	58
VIEWING CONTROLLER INFORMATION	58
MAKING CONTROLLER SETTINGS	59
VIEWING CONTROLLER STATISTICS	60
VIEWING THE FLASH IMAGE INFORMATION	60

UPDATING FIRMWARE ON A RAID SUBSYSTEM	61
VIEWING BATTERY INFORMATION	62
RECONDITIONING A BATTERY	62
BUZZER SETTINGS	63
SILENCING THE BUZZER	63
MANAGING ENCLOSURES	63
VIEWING ENCLOSURE TOPOLOGY	64
VIEWING THE ENCLOSURES SUMMARY	65
LOCATING AN ENCLOSURE	65
VIEWING ENCLOSURE INFORMATION	65
MAKING ENCLOSURE SETTINGS	65
VIEWING FRU VPD INFORMATION	66
VIEWING POWER SUPPLY STATUS	66
VIEWING COOLING UNIT STATUS	66
VIEWING TEMPERATURE SENSOR STATUS	66
VIEWING VOLTAGE SENSOR STATUS	66
MANAGING UPS UNITS	67
VIEWING A LIST OF UPS UNITS	67
MAKING UPS SETTINGS	67
VIEWING UPS INFORMATION	68

MANAGING NETWORK CONNECTIONS	68
MAKING VIRTUAL MANAGEMENT PORT SETTINGS	68
MAKING MAINTENANCE MODE SETTINGS	69

MANAGING USERS	69
VIEWING USER INFORMATION	69
CREATING A USER	70
MAKING USER SETTINGS	70
CHANGING USER PASSWORDS	70
DELETING A USER	71
SETTING USER EVENT SUBSCRIPTIONS	71
IMPORTING A USER DATABASE	71
EXPORTING A USER DATABASE	72

MANAGING BACKGROUND ACTIVITIES	72
VIEWING CURRENT BACKGROUND ACTIVITIES	73
VIEWING SCHEDULED BACKGROUND ACTIVITIES	73
ADDING A SCHEDULED BACKGROUND ACTIVITY	73
CHANGING A BACKGROUND ACTIVITY SCHEDULE	74
ENABLING OR DISABLING A SCHEDULED BACKGROUND ACTIVITY	74
DELETING A SCHEDULED BACKGROUND ACTIVITY	74
MEDIA PATROL	75
REDUNDANCY CHECK	75
INITIALIZATION	75
REBUILD	76
MIGRATION	76
PDM	76
TRANSITION	77
SYNCHRONIZATION	77

BATTERY RECONDITIONING	77
MANAGING STORAGE SERVICES	78
VIEWING A LIST OF SERVICES	78
EMAIL SERVICE	78
SLP SERVICE	79
WEBSERVER SERVICE	79
TELNET SERVICE	80
SSH SERVICE	80
SNMP SERVICE	81
NETSEND SERVICE	82

WORKING WITH THE EVENT VIEWER	84
VIEWING RUNTIME EVENTS	84
SAVING RUNTIME EVENTS	84
CLEARING RUNTIME EVENTS	84
VIEWING NVRAM EVENTS	85
SAVING NVRAM EVENTS	85
CLEARING NVRAM EVENTS	85

MONITORING PERFORMANCE	85
MONITORING I/O PERFORMANCE	85
MONITORING PSU WATTAGE	86

MANAGING PHYSICAL DRIVES	87
VIEWING A LIST OF PHYSICAL DRIVES	87
VIEWING PHYSICAL DRIVE INFORMATION	87
MAKING GLOBAL PHYSICAL DRIVE SETTINGS	88
MAKING INDIVIDUAL PHYSICAL DRIVE SETTINGS	88
VIEWING PHYSICAL DRIVE STATISTICS	88
VIEWING PHYSICAL DRIVE SMART LOG INFORMATION	89
SAVING THE PHYSICAL DRIVE SMART LOG	89
LOCATING A PHYSICAL DRIVE	89
FORCING A PHYSICAL DRIVE OFFLINE	90
CLEARING A STALE OR A PFA CONDITION	90
UPDATING FIRMWARE ON A PHYSICAL DRIVE	90

MANAGING DISK ARRAYS	91
VIEWING A LIST OF DISK ARRAYS	91
VIEWING DISK ARRAY INFORMATION	91
CREATING A DISK ARRAY MANUALLY	92
CREATING A DISK ARRAY WITH THE WIZARD	92
DELETING A DISK ARRAY	93
LOCATING A DISK ARRAY	93
MAKING DISK ARRAY SETTINGS	93
RUNNING MEDIA PATROL ON A DISK ARRAY	94
RUNNING PDM ON A DISK ARRAY	94
PREPARING A DISK ARRAY FOR TRANSPORT	95
REBUILDING A DISK ARRAY	95

MANAGING LOGICAL DRIVES	96
VIEWING A LIST OF LOGICAL DRIVES	96
VIEWING LOGICAL DRIVE INFORMATION	96
VIEWING LOGICAL DRIVE STATISTICS	97

VIEWING LOGICAL DRIVE CHECK TABLES	97
CREATING A LOGICAL DRIVE MANUALLY	97
DELETING A LOGICAL DRIVE	98
MAKING LOGICAL DRIVE SETTINGS	98
LOCATING A LOGICAL DRIVE	98
INITIALIZING A LOGICAL DRIVE	99
REDUNDANCY CHECK ON A LOGICAL DRIVE	99
MIGRATING A LOGICAL DRIVE'S RAID LEVEL	100
CREATING A LUN CLONE	100

MANAGING SPARE DRIVES	101
VIEWING A LIST OF SPARE DRIVES	101
VIEWING SPARE DRIVE INFORMATION	102
CREATING A SPARE DRIVE MANUALLY	102
DELETING A SPARE DRIVE	102
MAKING SPARE DRIVE SETTINGS	102
LOCATING A SPARE DRIVE	103
RUNNING SPARE CHECK	103
RUNNING A TRANSITION ON A SPARE DRIVE	103

MANAGING INITIATORS	104
VIEWING A LIST OF INITIATORS	104
ADDING AN FC INITIATOR	104
DELETING AN FC INITIATOR	104
ADDING AN iSCSI INITIATOR	105

MANAGING LUNs	105
VIEWING A LIST OF LUN MAPS	105
LUN MAPPING AND MASKING	105
ADDING A LUN MAP	106
EDITING A LUN MAP	106
DELETING A LUN MAP	107
ENABLING AND DISABLING LUN MASKING	107

MANAGING FIBRE CHANNEL CONNECTIONS	107
VIEWING FC NODE INFORMATION	107
VIEWING FC PORT INFORMATION	107
MAKING FC PORT SETTINGS	108
VIEWING FC PORT STATISTICS	108
VIEWING A LIST OF FC INITIATORS ON THE FABRIC	108
VIEWING A LIST OF FC LOGGED-IN DEVICES	109
VIEWING A LIST OF FC SFPs	109

MANAGING iSCSI CONNECTIONS	109
VIEWING A LIST OF iSCSI TARGETS	110
VIEWING iSCSI TARGET INFORMATION	110
MAKING iSCSI TARGET SETTINGS	111
VIEWING A LIST OF iSCSI PORTALS	111
VIEWING iSCSI PORTAL INFORMATION	111
ADDING iSCSI PORTALS	111
MAKING iSCSI PORTAL SETTINGS	112
DELETING iSCSI PORTALS	112
VIEWING A LIST OF iSCSI PORTS	112

VIEWING iSCSI PORT INFORMATION	113
MAKING iSCSI PORT SETTINGS	113
VIEWING A LIST OF iSCSI TRUNKS	113
ADDING iSCSI TRUNKS	113
MAKING iSCSI TRUNK SETTINGS	114
DELETING iSCSI TRUNKS	114
VIEWING A LIST OF iSCSI SESSIONS	114
VIEWING iSCSI SESSION INFORMATION	114
DELETING AN iSCSI SESSION	115
VIEWING iSCSI iSNS INFORMATION	115
MAKING iSCSI iSNS SETTINGS	115
VIEWING A LIST OF iSCSI CHAPs	116
ADDING iSCSI CHAPs	116
MAKING iSCSI CHAP SETTINGS	116
DELETING iSCSI CHAPs	116
PINGING A HOST OR SERVER ON THE iSCSI NETWORK	117

CHAPTER 5: MANAGEMENT WITH THE CLU 118

INITIAL CONNECTION	119
MAKING A SERIAL CONNECTION	119
MAKING A TELNET CONNECTION	119
MAKING A SSH CONNECTION	120
LOGGING INTO THE CLI	120
ACCESSING ONLINE HELP	121
EXITING THE CLU	121
LOGGING OUT OF THE CLI	121
LOGGING BACK INTO THE CLI AND CLU	121
MANAGING THE SUBSYSTEM (CLU)	122
MAKING SUBSYSTEM SETTINGS (CLU)	122
RUNNING MEDIA PATROL (CLU)	122
LOCKING OR UNLOCKING THE SUBSYSTEM (CLU)	122
SETTING SUBSYSTEM DATE AND TIME (CLU)	123
MAKING NTP SETTINGS (CLU)	123
SYNCHRONIZING WITH A NTP SERVER (CLU)	123

MANAGING THE RAID CONTROLLERS (CLU)	124
VIEWING CONTROLLER INFORMATION (CLU)	124
CLEARING AN ORPHAN WATERMARK (CLU)	124
MAKING CONTROLLER SETTINGS (CLU)	124
LOCATING THE CONTROLLER (CLU)	125

MANAGING THE ENCLOSURE (CLU)	126
VIEWING THE ENCLOSURES SUMMARY (CLU)	126
VIEWING ENCLOSURE INFORMATION (CLU)	126
MAKING ENCLOSURE SETTINGS (CLU)	126
VIEWING FRU VPD INFORMATION (CLU)	127
VIEWING POWER SUPPLY STATUS (CLU)	127
TO VIEW THE STATUS OF THE POWER SUPPLIES:	127
LOCATING A POWER SUPPLY (CLU)	127
VIEWING COOLING UNIT STATUS (CLU)	127

VIEWING TEMPERATURE SENSOR STATUS (CLU)	127
VIEWING VOLTAGE SENSOR STATUS (CLU)	127
VIEWING BATTERY INFORMATION (CLU)	128
RECONDITIONING A BATTERY (CLU)	128
LOCATING AN ENCLOSURE (CLU)	128
VIEWING ENCLOSURE TOPOLOGY (CLU)	129
PHYSICAL DRIVE MANAGEMENT (CLU)	129
VIEWING A LIST OF PHYSICAL DRIVES (CLU)	129
MAKING GLOBAL PHYSICAL DRIVE SETTINGS (CLU)	129
VIEWING PHYSICAL DRIVE INFORMATION (CLU)	130
VIEWING PHYSICAL DRIVE STATISTICS (CLU)	130
SETTING AN ALIAS (CLU)	130
CLEARING STALE AND PFA CONDITIONS (CLU)	130
FORCING A PHYSICAL DRIVE OFFLINE (CLU)	131
LOCATING A PHYSICAL DRIVE (CLU)	131
MANAGING DISK ARRAYS (CLU)	132
VIEWING A LIST OF DISK ARRAYS (CLU)	132
CREATING A DISK ARRAY (CLU)	132
CREATING A DISK ARRAY – AUTOMATIC (CLU)	132
CREATING A DISK ARRAY – EXPRESS (CLU)	132
CREATING A DISK ARRAY – ADVANCED (CLU)	133
DELETING A DISK ARRAY (CLU)	134
MAKING DISK ARRAY SETTINGS (CLU)	134
VIEWING DISK ARRAY INFORMATION (CLU)	134
ACCEPTING AN INCOMPLETE ARRAY (CLU)	135
ENABLING MEDIA PATROL, PDM, AND POWER MANAGEMENT ON A DISK ARRAY (CLU)	135
PREPARING THE DISK ARRAY FOR TRANSPORT (CLU)	135
REBUILDING A DISK ARRAY (CLU)	135
RUNNING MEDIA PATROL ON A DISK ARRAY (CLU)	135
RUNNING PDM ON A DISK ARRAY (CLU)	136
RUNNING TRANSITION ON A DISK ARRAY (CLU)	136
LOCATING A DISK ARRAY (CLU)	136
MANAGING SPARE DRIVES (CLU)	137
VIEWING A LIST OF SPARE DRIVES (CLU)	137
CREATING A SPARE DRIVE (CLU)	137
MAKING SPARE DRIVE SETTINGS (CLU)	137
RUNNING SPARE CHECK (CLU)	138
DELETING A SPARE DRIVE (CLU)	138
To delete a spare drive:	138
MANAGING LOGICAL DRIVES (CLU)	138
CREATING A LOGICAL DRIVE (CLU)	138
DELETING A LOGICAL DRIVE (CLU)	139
VIEWING LOGICAL DRIVE INFORMATION (CLU)	139
VIEWING LOGICAL DRIVE STATISTICS (CLU)	139
VIEWING THE LOGICAL DRIVE CHECK TABLE (CLU)	139
MAKING LOGICAL DRIVE SETTINGS (CLU)	140
INITIALIZING A LOGICAL DRIVE (CLU)	140
RUNNING REDUNDANCY CHECK (CLU)	140

LOCATING A LOGICAL DRIVE (CLU)	140
MIGRATING A LOGICAL DRIVE (CLU)	141
CREATING A LUN CLONE (CLU)	141
MANAGING THE NETWORK CONNECTION (CLU)	142
MAKING VIRTUAL MANAGEMENT PORT SETTINGS (CLU)	142
MAKING MAINTENANCE MODE SETTINGS (CLU)	143
MANAGING FIBRE CHANNEL CONNECTIONS (CLU)	143
VIEWING NODE INFORMATION (CLU)	143
VIEWING FIBRE CHANNEL PORT INFORMATION	143
VIEWING FIBRE CHANNEL LOGGED-IN DEVICES (CLU)	144
MAKING FIBRE CHANNEL PORT SETTINGS (CLU)	144
VIEWING FIBRE CHANNEL PORT STATISTICS (CLU)	144
VIEWING SFP INFORMATION (CLU)	144
VIEWING FIBRE CHANNEL PORT STATISTICS (CLU)	145
VIEWING FIBRE CHANNEL INITIATORS (CLU)	145
MANAGING iSCSI CONNECTIONS (CLU)	146
VIEWING A LIST OF iSCSI TARGETS (CLU)	146
VIEWING iSCSI TARGET INFORMATION (CLU)	147
MAKING iSCSI TARGET SETTINGS (CLU)	147
VIEWING A LIST OF iSCSI PORTS (CLU)	148
VIEWING iSCSI PORT INFORMATION (CLU)	148
MAKING iSCSI PORT SETTINGS (CLU)	148
VIEWING A LIST OF iSCSI PORTALS (CLU)	148
VIEWING iSCSI PORTAL INFORMATION (CLU)	149
ADDING iSCSI PORTALS (CLU)	149
MAKING iSCSI PORTAL SETTINGS (CLU)	149
DELETING iSCSI PORTALS (CLU)	150
VIEWING A LIST OF iSCSI SESSIONS (CLU)	150
DELETING AN iSCSI SESSION (CLU)	150
VIEWING iSCSI SESSION INFORMATION (CLU)	151
VIEWING iSCSI iSNS INFORMATION (CLU)	151
MAKING iSCSI iSNS SETTINGS (CLU)	151
VIEWING A LIST OF iSCSI CHAPs (CLU)	152
ADDING iSCSI CHAPs (CLU)	152
MAKING iSCSI CHAP SETTINGS (CLU)	152
DELETING iSCSI CHAPs (CLU)	152
PINGING A HOST OR SERVER ON THE iSCSI NETWORK (CLU)	152
VIEWING A LIST OF iSCSI TRUNKS (CLU)	153
ADDING iSCSI TRUNKS (CLU)	153
MAKING iSCSI TRUNK SETTINGS (CLU)	153
DELETING iSCSI TRUNKS (CLU)	153
MANAGING BACKGROUND ACTIVITY (CLU)	154
VIEWING CURRENT BACKGROUND ACTIVITIES (CLU)	154
MAKING BACKGROUND ACTIVITY SETTINGS (CLU)	154
WORKING WITH THE EVENT VIEWER (CLU)	155
VIEWING RUNTIME EVENTS (CLU)	155
CLEARING RUNTIME EVENTS (CLU)	155
VIEWING NVRAM EVENTS (CLU)	155

CLEARING NVRAM EVENTS (CLU)	155
WORKING WITH LUN MAPPING (CLU)	156
ENABLING LUN MAPPING (CLU)	156
VIEWING A LIST OF INITIATORS (CLU)	156
ADDING AN INITIATOR (CLU)	156
DELETING AN INITIATOR (CLU)	157
VIEWING A LIST OF LUN MAPS (CLU)	157
ADDING A LUN MAP (CLU)	157
EDITING A LUN MAP (CLU)	158
DELETING A LUN MAP (CLU)	158
CHANGING THE ACTIVE LUN MAPPING TYPE (CLU)	158
MANAGING UPS UNITS (CLU)	159
VIEWING A LIST OF UPS UNITS (CLU)	159
MAKING UPS SETTINGS (CLU)	159
VIEWING UPS INFORMATION (CLU)	160
MANAGING USERS (CLU)	160
VIEWING USER INFORMATION (CLU)	160
CREATING A USER (CLU)	160
CHANGING ANOTHER USER’S SETTINGS (CLU)	161
CHANGING YOUR OWN USER SETTINGS (CLU)	161
CHANGING ANOTHER USER’S PASSWORD (CLU)	161
DELETING A USER (CLU)	162
WORKING WITH SOFTWARE MANAGEMENT (CLU)	163
MAKING EMAIL SETTINGS (CLU)	163
MAKING SLP SETTINGS (CLU)	163
MAKING WEB SERVER SETTINGS (CLU)	163
MAKING TELNET SETTINGS (CLU)	164
MAKING SSH SETTINGS (CLU)	164
MAKING SNMP SETTINGS (CLU)	164
MANAGING SNMP TRAP SINKS (CLU)	164
MAKING NETSEND SETTINGS (CLU)	165
MANAGING NETSEND RECIPIENTS (CLU)	165
FLASHING THROUGH TFTP	166
VIEWING FLASH IMAGE INFORMATION (CLU)	166
CLEARING STATISTICS (CLU)	167
RESTORING FACTORY DEFAULTS (CLU)	167
SHUTTING DOWN THE SUBSYSTEM (CLU)	168
SHUTTING DOWN THE VESS R2600 – TELNET CONNECTION	168
SHUTTING DOWN THE VESS R2600 – SSH CONNECTION	168
SHUTTING DOWN THE VESS R2600 – SERIAL CONNECTION	168
STARTING UP AFTER SHUTDOWN	169
STARTING UP THE VESS R2600 – TELNET CONNECTION	169
STARTING UP THE VESS R2600 – SSH CONNECTION	169
STARTING UP THE VESS R2600 – SERIAL CONNECTION	169
RESTARTING THE SUBSYSTEM	170
RESTARTING VESS R2600 – TELNET CONNECTION	170
RESTARTING VESS R2600 – SSH CONNECTION	171

RESTARTING VESS R2600 — SERIAL CONNECTION	171
BUZZER	171
MAKING BUZZER SETTINGS	171
SILENCING THE BUZZER	171
CHAPTER 6: MAINTENANCE	172
UPDATING THE SUBSYSTEM FIRMWARE	172
UPDATING WITH WebPAM PROe	172
UPDATING WITH THE CLU	173
UPDATING WITH USB SUPPORT	174
UPDATING PHYSICAL DRIVE FIRMWARE	175
WebPAM PROe	175
RESTARTING A SUBSYSTEM	176
REPLACING A POWER SUPPLY	177
REMOVING THE OLD POWER SUPPLY	177
INSTALLING A NEW POWER SUPPLY	177
REPLACING A COOLING UNIT	178
REPLACING A CACHE BACKUP BATTERY	179
REPLACING A RAID CONTROLLER — DUAL CONTROLLERS	180
REMOVING THE OLD CONTROLLER	180
INSTALLING THE NEW CONTROLLER	181
REPLACING A RAID CONTROLLER — SINGLE CONTROLLER	181
REMOVING THE OLD CONTROLLER	182
INSTALLING THE NEW CONTROLLER	182
CHAPTER 7: TECHNOLOGY BACKGROUND	183
DISK ARRAYS	183
MEDIA PATROL	183
PDM	184
LOGICAL DRIVES	184
RAID LEVELS	184
RAID LEVEL MIGRATION	189
STRIPE SIZE	192
SECTOR SIZE	192
PREFERRED CONTROLLER ID	192
INITIALIZATION	192
PARTITION AND FORMAT	192
SPARE DRIVES	193
DEFINITION	193
OPTIONS	193
REQUIREMENTS	193
TRANSITION	193
RAID CONTROLLERS	194
LUN AFFINITY	194
ALUA	194
CACHE POLICY	195
PREFERRED CONTROLLER ID	196
POWER SAVING	196

CAPACITY COERCION	196
iSCSI MANAGEMENT	197
iSCSI ON A VLAN	197
INITIATOR	198
TARGET	198
PORT	199
TRUNK	199
SESSION	199
iSNS	199
CHAP	199
PING	199
INTERNET PROTOCOLS	200
CHAPTER 8: TROUBLESHOOTING	201
VESS R2600 IS BEEPING	201
SILENCING THE BUZZER	201
LEDs DISPLAY AMBER OR RED	202
FRONT PANEL LEDs	202
DRIVE CARRIER LEDs	203
BACK PANEL LEDs	203
CHECKING COMPONENT INSTALLATION	206
CLU REPORTS A PROBLEM	206
VIEWING RUNTIME EVENTS	206
VIEWING NVRAM EVENTS	207
CHECKING A REPORTED COMPONENT	207
WebPAM PROe REPORTS A PROBLEM	208
HEADER	208
DASHBOARD TAB	208
DEVICE TAB	208
STORAGE TAB	208
ADMINISTRATION TAB	208
USB SUPPORT REPORTS A PROBLEM	209
ENCLOSURE PROBLEMS	209
DIAGNOSING AN ENCLOSURE PROBLEM	209
OVERHEATING	210
POWER SUPPLIES	210
BATTERIES	210
RAID CONTROLLER PROBLEMS	211
MAINTENANCE MODE	211
FINDING AND CORRECTING THE CAUSE OF THE PROBLEM	211
TAKING A RAID CONTROLLER OUT OF MAINTENANCE MODE	211
UNSAVED DATA IN THE CONTROLLER CACHE	212
PHYSICAL DRIVE PROBLEMS	213
DISK ARRAY AND LOGICAL DRIVE PROBLEMS	213
DISK ARRAY DEGRADED / LOGICAL DRIVE CRITICAL	213
DISK ARRAY OFFLINE / LOGICAL DRIVE OFFLINE	214
REPAIRING AN OFFLINE DISK ARRAY OR LOGICAL DRIVE	214

REBUILDING A DISK ARRAY	215
INCOMPLETE ARRAY	215
CONNECTION PROBLEMS	216
SERIAL CONNECTIONS	216
NETWORK CONNECTIONS	216
FIBRE CHANNEL CONNECTIONS	217
SAS CONNECTIONS	217
BROWSER DOES NOT CONNECT TO WebPAM PROe	217
POWER CYCLING THE SUBSYSTEM	218
EVENT NOTIFICATION RESPONSE	218
CHAPTER 9: SUPPORT	225
FREQUENTLY ASKED QUESTIONS	225
PHYSICAL DRIVES	225
STARTUP	226
LOGGING IN	226
CONNECTION	226
TIMEOUTS	226
EMAIL MESSAGES	226
USER MANAGEMENT	226
LOCK	226
CREATING A DISK ARRAY OR LOGICAL DRIVE	226
DISK ARRAY DEGRADED	227
DELETING A DISK ARRAY OR LOGICAL DRIVE	227
REBUILDING A DISK ARRAY	227
MEDIA PATROL AND PDM	227
CONTACTING TECHNICAL SUPPORT	228
APPENDIX A: USEFUL INFORMATION	232
SNMP MIB FILES	232
ADDING A SECOND RAID CONTROLLER	233
INSTALLING A SECOND RAID CONTROLLER	233
RAID CONTROLLER IN MAINTENANCE MODE	233
NEW SETTINGS FOR DUAL CONTROLLERS	233
DUAL CONTROLLERS AND SATA DRIVES	234

Appendix B: Multipathing on Windows 235

Before You Begin	235
Installing PerfectPath	236
Verifying Installation	236
Start Menu	236
Services List	236
Device Manager	236
Running Perfect Path View	237
Starting PerfectPath View	237
Quitting PerfectPath View	237
Monitoring Your LUNs and Paths	238
Viewing LUN Properties	238
Viewing Path Properties	238
Viewing LUN Performance Statistics	238
Viewing Path Performance Statistics	239
Viewing Events	239
Clearing Path Statistics	239
Features and Settings	240
Automatic Load Balancing for Failover Policy	240
Load Balance Policy	240
Path Verification	241
PDO Removal	241
Performance Tab Refresh Rate	242
Round Robin Count	242
Refreshing the Objects	242
Viewing System Information	242
Saving System Information	243
Troubleshooting	243
Updating PerfectPath	244
Repairing PerfectPath	244
Removing PerfectPath	245
Preferred Method	245
Alternate Method 1	245
Alternate Method 2	245

Appendix C: Multipathing on Linux 246

Before You Begin	246
Check Initial Setup	246
Task 1: Meeting Package Requirements	247
Installing Packages	247
Verifying Packages – RedHat	247
Verifying Packages – SuSE	247
Task 2: Preparing the Configuration File	248
RedHat Systems	248
SuSE Systems	248
Editing a Configuration File	248
Task 3: Making Initial Host Settings	249

Setting the Daemon to Run	249
Verifying the Modules are Loaded	250
Verifying the Daemon is Running	250
Task 4: Create and Configure Devices	250
Task 5: Setting-up ALUA	251
RedHat Linux RHEL 5.4	251
RedHat Linux RHEL 5.5	251
SuSE Linux SLES 10 SP3	251
SuSE Linux SLES 11	252
SuSE Linux SLES 11 SP1	252
RPM Packages and Documents for Linux	253
Linux MPIO: Known Issues	254
Sample multipath.conf File	255

CHAPTER 1: INTRODUCTION

This chapter covers the following topics:

- "About This Manual" on page 2
- "Vess R2600 Overview" on page 3
- "New Features in the Vess R2600" on page 3
- "Architectural Description" on page 4
- "Features" on page 6
- "General Specifications" on page 9
- "Warranty and Support" on page 11

ABOUT THIS MANUAL





This **Product Manual** describes how to setup, use, and maintain the Vess R2000 Series and Vess J2600 external disk array subsystems. The manual is organized into chapters as follows:

- "Chapter 1: Introduction" on page 1, this chapter provides a general overview of the available devices in the Vess R2000 Series and Vess J2600.
- "Chapter 2: Hardware Installation" on page 12 describes the steps necessary for installing subsystem hardware including installing hard disks and placing the device into a rack system.
- "Chapter 3: Setup" on page 35 describes setting up a serial connection and the basics of how to use the built-in command-line interface (CLI), the built-in command-line utility (CLU), and the embedded Web-based Promise Array Management – Professional (WebPAM PROe) software.
- "Chapter 4: Manage with WebPAM PROe" on page 50 provides a more detailed description of the various menus used for managing the Vess R2600 and connected Vess J2600 expansion devices.
- "Chapter 5: Management with the CLU" on page 118 describes using the CLU and CLI (access to the CLU is done through the CLI) to manage the Vess R2600 through the network or via serial connection.
- "Chapter 6: Maintenance" on page 172 describes how to replace hardware components including RAID controllers, power supplies, and cooling units; how to update firmware for subsystems and physical drives.
- "Chapter 7: Technology Background" on page 183 provides a description of the technologies and concepts that underlie networked RAID storage systems generally and the Vess R2600 and Vess J2600 subsystems in particular.
- "Chapter 8: Troubleshooting" on page 201 describes what to do in response to specific problems that might be encountered over the lifetime operation of the Vess R2600 and Vess J2600 subsystems. Included in the chapter are descriptions of the various types of alerts and notices delivered though the management interfaces (WebPAM PROe, CLU, CLI) or hardware (LEDs and audible signals).
- "Chapter 9: Support" on page 225 includes Frequently Asked Questions (FAQ), how to contact technical support, how to return a system for repair, and warranty information.
- "Appendix A: Useful Information" on page 232
- "Appendix B: Multipathing on Windows" on page 235
- "Appendix C: Multipathing on Linux" on page 246

This manual includes a full table of contents, index, chapter task lists and numerous cross-references to help you find the specific information you are looking for.

The terms “Vess R2600” or “subsystem” are used in examples or descriptions throughout this manual to refer to any of the available Vess R2000 Series models. The terms “unit” or “device” can refer to any Vess R2000 Series or Vess J2600 model.

Also included are four levels of notices:

	Warning A Warning notifies you of probable equipment damage or loss of data, or the possibility of physical injury, and how to avoid them.
	Caution A Caution informs you of possible equipment damage or loss of data and how to avoid them.
	Important An Important message calls attention to an essential step or point required to complete a task, including things often missed.
	Note A Note provides helpful information such as hints or alternative ways of doing a task.

Vess R2600 Overview

All PROMISE Vess R2600 and Vess J2600 2000 models support use of 6 Gb/s SAS and SATA disks.

The Vess 2600 models are coupled with high speed 8 Gb/s Fibre Channel host connectivity and 1 Gb/s iSCSI host connectivity.

The Vess R2600 model is coupled with high speed 1 Gb/s iSCSI host connectivity.

Performance

The PROMISE Vess R2600 is built using a 64bit 6-core processor per RAID controller, and support for 6 Gb/s SAS and SATA hard disk drives and solid state drives (SR1). Dual active-active controller modules with cache mirroring over a PCIe Gen 2 link allow for redundant data paths to ensure data availability while dual power supply/cooling units minimize downtime and any disruption to business continuity.

GreenRAID

PROMISE cares about the environment. Vess R2600 products utilize environmentally friendly production methods and materials and are designed with high-efficiency in mind. Powered by 80Plus certified power supplies that offer up to 85% power efficiency, all Vess R2600 devices improve total cost of ownership by conserving power, reducing heat output and improving cooling costs. Additionally, the PROMISE as GreenRAID story includes advanced power management support for hard disk drives providing up to 65% energy savings.

Service and Support

Every Vess R2600 subsystem is backed by the PROMISE Three-Year limited warranty with 24-hour, 7-day English language telephone and e-mail support. In addition to our industry leading warranty, PROMISE offers extended warranty and onsite parts replacement options with service levels with response times as low four hours.

New Features in the Vess R2600

PERFECTREBUILD™

The PerfectRebuild™ feature is an innovative approach to rebuilding a RAID array in order to significantly reduce the amount of time needed for completion. This frees up CPU resources more quickly to be available for I/O and other demands. PerfectRebuild™ ignores any portion of the logical drive where no write changes have occurred, focusing only on the parts that have changed. The conventional approach has been to rebuild the entire logical drive, even sections with no write changes. This reduction in the total time needed for a rebuild is especially significant for very large drives.

ADVANCED BATTERY FLASH BACKUP

Use the optional Backup Battery Unit (BBU) and Flash modules for maximum data protection in the event of power loss. High performance RAID arrays operating in write-back mode present a write loss risk in the event of a power loss. Since data is considered committed as soon as the controller has received it in cache memory, it will be lost if power is interrupted. The conventional solution is to use a BBU to maintain the RAID cache for 72 hours. A better solution is to use power from the BBU to write the content of the RAID controller write cache to non-volatile flash memory in order to extend the period of cache protection beyond the standard 72 hours provided by a typical BBU.

HARD DISK CACHE PROTECTION

HD Cache Protection is another capability available for use with high performance RAID arrays that employ the optional BBU and USB Flash modules. HD Cache Protection uses the BBU and USB Flash modules to protect against loss of HDD cache data, eliminating the need to add an expensive universal power supply (UPS). The HDD cache is written to flash memory, again, using power from the BBU in order to protect the HDD cache beyond the standard 72 hours. (SR1)

ARCHITECTURAL DESCRIPTION

The Vess R2600 subsystems are suitable for Direct Attached Storage (DAS), Storage Area Network (SAN), and Expanded Storage with the Vess J2600.

Vess R2000 SERIES MODEL LINE-UP

Model	Controller Units	Interface	Number of Drives	Power Supplies	Controller Fans
R2600fiD	2	FC/iSCSI	16	3	2
R2600iD	2	iSCSI	16	3	2
R2600fiS	1	FC/iSCSI	16	3	2
R2600iS	1	iSCSI	16	3	2

Vess J2000 SERIES MODEL LINE-UP


Model	Controller Units	Interface	Number of Drives	Power Supplies	Controller Fans
J2600sD	2	SAS	16	3	2
J2600sS	1	SAS	16	3	2

Figure 1 - 1: Vess R2600 front view




- 16-bay SATA/SAS

Figure 1 - 2: Vess R2600fiD back view




- Two 8Gb/s FC ports per controller
- Four 1Gb/s iSCSI ports per controller
- One JBOD Expansion port per controller
- Three PSU
- Two Cooling Units (each Cooling Unit for Vess R2600 models include two fans and a BBU)

Figure 1 - 3: Vess R2600iD (pictured with optional PSU 4) back view

- Four 1Gb/s iSCSI ports per controller
- One JBOD Expansion port per controller
- Three PSU
- Two Cooling Units (each Cooling Unit for Vess R2600 models include two fans and a BBU)

Figure 1 - 4: Vess J2600 front view

- 16-bay SATA/SAS

Figure 1 - 5: Vess J2600sD back view

- SAS port per controller
- Three PSU
- Two Cooling Units (each Cooling Unit for Vess J2600 models include two fans)

FEATURES

CONTROLLER MODULE FEATURES

DRIVE SUPPORT

- 3.5-inch and 2.5-inch form factor
- Hard disk drives (HDDs) and Solid State drives (SSDs)
- SAS, 6 Gb/s and 3 Gb/s
- SATA, 6 Gb/s and 3 Gb/s
- Supports any mix of SAS and SATA drives simultaneously in the same enclosure

For a list of supported drives, go to PROMISE support:

<http://www.promise.com/support/>

SATA physical drives require a SAS-to-SATA adapter, available from PROMISE.

EXTERNAL I/O PORTS PER CONTROLLER

- Vess R2600fi: Two 8 Gb/s Fibre Channel ports compatible with 4 Gb/s and 2 Gb/s and four 1 Gb/s iSCSI ports
- Vess R2600i: Four 1 Gb/s iSCSI ports
- All models: One external SAS port with an SFF-8088 SAS connector, supports up to 4 cascading JBOD expansion units (FCS), up to 6 cascading JBOD expansion units (SR1)

DATA CACHE

- 2 GB data cache per controller.
A portion of the data cache is shared with the controller firmware
- Protected with hot-swappable battery backup unit (BBU)



Note

A maximum of 30 logical drives (LD) are supported.

OPERATIONAL FEATURES

RAID LEVEL SUPPORT

- 0, 1, 1E, 3, 5, 6, 10, 30, 50, and 60

RAID STRIPE SIZE SUPPORT

- 64K, 128K, 256K, 512K, and 1MB

HOT SPARE DRIVES

- Global
- Dedicated
- Revertible option

MAXIMUM LUNs SUPPORTED

- Total LUNs: 256
- Per Array: 32

ADVANCED STORAGE FEATURES

- Perfect Rebuild
- Advanced Battery Flash Backup
- HDD Cache Backup (SR3)
- Online LUN Clone (SR1)
- Advanced Cache Mirroring over PCIe Gen2
- Simple, drag-and-drop LUN Masking and Mapping
- Asymmetric LUN Unit Access (ALUA)
- Volume Copy
- PerfectFlash - Non-Disruptive Software Update
- I/O performance & power monitoring tools
- Guaranteed Latency Technology (an advanced OEM feature)
- USB Service Log

BACKGROUND ACTIVITIES

- Media Patrol
- Background Synchronization
- Foreground Initialization
- Rebuild
- Redundancy Check
- Disk SMART Polling
- Online Capacity Expansion (OCE)
- RAID Level Migration (RLM)
- UPS Monitoring
- Feature rich task scheduler for background activities

PERFECTRAID FEATURES

- Predictive Data Migration (PDM)
- Intelligent Bad Sector Remapping
- SMART Error Handling
- NVRAM Error Logging
- Disk Slot Power Control
- Read/Write Check Table
- Write Hole Table

GREENRAID FEATURES

- Four levels of advanced power management disk drive (MAID) support
- Efficient 80Plus Bronze Certified power supplies

SYSTEM MANAGEMENT

MANAGEMENT INTERFACES



Note

There are two options to provide the physical connection for system management for the Vess R2600 models, an RJ-11 serial port or an RJ-45 Ethernet network port. An RJ-11-to-DB9 adapter is shipped with each Vess R2600 model.

- Browser-based management with WebPAM PROe over Ethernet
- Command Line Interface (CLI) over Serial Port, Ethernet via Telnet, or SSH
- Command Line Utility (CLU) over Serial Port, Ethernet via Telnet, or SSH
- Third Party Management Support via SNMP

SUPPORTED OPERATING SYSTEMS

Operating systems run on the Host PC, from which you monitor and manage the Vess R2600 subsystem.

Supported Operating Environments		
Core Platform	Type	Notes
Microsoft		
Windows Server 2003 Enterprise Edition R2 with SP2	x86 / x64	
Windows Server 2008 Datacenter Edition with SP2	x86/ x64	
Windows Server 2008 R2 SP1 Datacenter Edition	x64	
RedHat		
Enterprise Linux 5.7	x86 / x64	
Enterprise Linux 6.3	x86 / x64	
SuSE		
Linux Enterprise Server 10 + SP4	x86 / x64	SUSE LINUX Enterprise Server 10 + SP4-32bit can't support dual controller boot from SAN SUSE LINUX Enterprise Server 11 + SP2-32bit can't support dual controller boot from SAN
Linux Enterprise Server 11 + SP2	x86 / x64	

SUPPORTED BROWSERS

Browsers run on the host PC or server, from which you monitor and manage the Vess R2600 subsystem using WebPAM PROe. The browsers listed here meet the minimum version requirements for browser compatibility:

- Mozilla Firefox 14.0.1
- Google Chrome 20.0.1132.57 m
- Internet Explorer 7 (Version: 7.0.5730.13)
- Internet Explorer 8 (Version: 8.0.6001.18702)
- Internet Explorer 9 (Version:)
- Safari 5.1.7 for Windows
- Safari 5.1.7 for MAC
- Mozilla Firefox for Linux 3.6.13

For the latest list of supported browsers, go to PROMISE support:
<http://www.promise.com/support/>

GENERAL SPECIFICATIONS

POWER SUPPLIES

- Hot swap and redundant with PFC, 100-240 Vac auto-ranging, 50 -60 Hz, N+1 design. Meets 80Plus bronze.

VOLTAGE

- 100-240 VAC
- Auto-Ranging

CURRENT (MAXIMUM)

- 4 A @ 115 VAC
- 2 A @ 230 VAC

POWER CONVERSION EFFICIENCY

- >80% @ 110V (>20% load)
- >80% @ 240V (>20% load)

OPERATING ENVIRONMENT

TEMPERATURE RANGE

- Operational: 5° to 35°C (41° to 95°F)
- Non-Operational: -40° to 60°C (-40° to 140°F)

HUMIDITY RANGE

- Operational: 10% to 90% (Non-Condensing)
- Non-Operational: 5% to 95% (Non-Condensing)

NOISE, SHOCK, AND VIBRATION

ACOUSTIC NOISE LEVELS

- Typical: 55 dB
- Maximum: 75 dB

SHOCK

- Operational: 5G, 11 ms duration
- Non-Operational: 30G, 11ms duration

VIBRATION

- Operational: 0.2G, sinewave, 0.5 oct/min, 5 to 500 Hz
- Non-Operational: 1G, 5 to 500 Hz

DIMENSIONS

(Height, Width, Depth)

- 131 x 447 x 507 mm (5.2 x 17.6 x 19.96 in)

WEIGHT

- 31.3 kg / 69 lbs (with drives installed)
- 20.1 kg / 44.3 lbs (without drives installed)

SAFETY AND ENVIRONMENTAL

EMI / RFI STATEMENTS

BSMI

CE

Warning: This is a class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

FCC

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

GOST-R

KCC

VCCI

ENVIRONMENTAL STANDARDS

- RoHS
- GreenPC
- WEEE

WARNINGS AND CAUTIONS

Warnings and Cautions are placed in this **Product Manual** beside the user actions to which they apply.

You can find these warnings and cautions under:

- "Unpacking the Vess R2600" on page 13
- "Mounting Vess R2600 in a Rack" on page 14
- "Installing Physical Drives" on page 16
- "Logging into WebPAM PROe" on page 44
- "Restoring Factory Default Settings" on page 55
- "Importing a Configuration Script" on page 57
- "Exporting a Configuration Script" on page 57
- "Silencing the Buzzer" on page 63
- "Making Virtual Management Port Settings" on page 68
- "Importing a User Database" on page 71
- "Changing a Background Activity Schedule" on page 74
- "Enabling or Disabling a Scheduled Background Activity" on page 74
- "Battery Reconditioning" on page 77
- "Forcing a Physical Drive Offline" on page 90
- "Deleting a Disk Array" on page 93
- "Deleting a Logical Drive" on page 98
- "Initializing a Logical Drive (CLU)" on page 140
- "Deleting an FC Initiator" on page 104
- "Updating with WebPAM PROe" on page 172
- "Updating with the CLU" on page 173
- "Updating with USB Support" on page 174
- "Updating Physical Drive Firmware" on page 175
- "Replacing a Cache Backup Battery" on page 179
- "Replacing a RAID Controller – Single Controller" on page 181
- "Initialization" on page 192
- "USB Support Reports a Problem" on page 209
- "Unsaved Data in the Controller Cache" on page 212

WARRANTY AND SUPPORT

WARRANTY

- Three year complete system limited warranty
- Battery Backup Unit has a one year limited warranty
- Optional 2-year extended warranty
- Optional onsite parts replacement program

Promise Technology, Inc. ("Promise") warrants that for three (3) years from the time of the delivery of the product to the original end user except for one (1) year warranty on the battery backup unit:

- a) the product will conform to Promise’s specifications;
- b) the product will be free from defects in material and workmanship under normal use and service.

This warranty:

- a) applies only to products which are new and in cartons on the date of purchase;
- b) is not transferable;
- c) is valid only when accompanied by a copy of the original purchase invoice;
- d) is not valid on spare parts.

This warranty shall not apply to defects resulting from:

- a) improper or inadequate maintenance, or unauthorized modification(s), performed by the end user;
- b) operation outside the environmental specifications for the product;
- c) accident, misuse, negligence, misapplication, abuse, natural or personal disaster, or maintenance by anyone other than a Promise or a Promise authorized service center.

SUPPORT

- 24 hour, 7 days a week, 365 days a year e-mail and phone support (English only)
- 24 hour, 7 days a week, 365 days a year access to PROMISE support site
- Firmware and compatibility lists

CHAPTER 2: HARDWARE INSTALLATION

This chapter covers the following topics:

- "Unpacking the Vess R2600" on page 13
- "Mounting Vess R2600 in a Rack" on page 14
- "Installing Physical Drives" on page 16
- "Making Management and Data Connections" on page 18
- "Making Serial Cable Connections" on page 29
- "Connecting the Power" on page 30
- "LED Behavior" on page 30

UNPACKING THE VESS R2600

The Vess R2600 box contains the following items:

- Vess R2600 Unit (PTVR2K)
- *Quick Start Guide* printed
- RJ11-to-DB9 serial data cable
- Screws for disk drives
(70 pieces for 16-bay)
- 1.5m (4.9 ft) Power cords (3 cords for 3 PSU installed, 4 cords for 4 PSU installed)
- CD with SNMP files, **Product Manual** and **Quick Start Guide** in PDF format
- Sliding rail assembly for rack mounting

The Vess J2600 box contains the following items:

- Vess J2600 Unit (PTVJ2K)
- Quick Start Guide printed
- RJ11-to-DB9 serial data cable
- Screws for disk drives
(70 pieces for 16-bay)
- Sliding rail assembly for rack mounting
- 1.5m (4.9 ft) Power cords (3 cords for 3 PSU installed, 4 cords for 4 PSU installed)
- SAS cable (1 cable J2600sS / 2 cables J2600sD)
- CD with SNMP files, **Product Manual** and **Quick Start Guide** in PDF format



Warning

The electronic components within the Vess R2600 enclosure are sensitive to damage from Electro-Static Discharge (ESD). Observe appropriate precautions at all times when handling the Vess R2600 or its subassemblies.



Cautions

There is a risk of explosion if battery is replaced by an incorrect type.
Dispose of used batteries according to the instructions.

MOUNTING VESS R2600 IN A RACK



Cautions

- Do not populate any unit with hard drives until it has been securely installed in the rack.
- At least two persons are required to safely lift, place, and attach the Vess R2600 or Vess J2600 unit into a rack system.
- Do not lift or move the Vess R2600 or Vess J2600 unit by the handles, power supplies or the controller units. Hold the subsystem itself.
- Do not install the Vess R2600 or Vess J2600 unit into a rack without rails to support the subsystem.
- Only a qualified technician who is familiar with the installation procedure should mount and install the Vess R2600 or Vess J2600 unit.
- Mount the rails to the rack using the appropriate screws and flange nuts, fully tightened, at each end of the rail.
- Do not load the rails unless they are installed with screws as instructed.
- The rails available for the PROMISE Vess R2600 or Vess J2600 unit are designed to safely support that PROMISE Vess R2600 or Vess J2600 unit when properly installed. Additional loading on the rails is at the customer's risk.
- PROMISE Technology, Inc. cannot guarantee that the mounting rails will support your PROMISE Vess R2600 or Vess J2600 unit unless you install them as instructed.



Note


In order to reduce the weight of the Vess R2600 enclosure, you can remove the power supplies. If your Vess R2600 shipped with physical drives installed, you can remove all of the drive carriers as well.

Mounting rails are included with the Vess R2600.

To install the Vess R2600 subsystem into your rack:

1. Attach the mounting rail assemblies to the rack posts, using screws and nuts from your rack system.
 - The rail halves are riveted together and use no adjustment screws.
 - The front-left and front-right mounting rail ends are labeled.

Figure 2 - 1: Installing the rails onto your rack



- All rail ends, front and rear, attach at the outside of the rack posts.
- The guide pins at the rail ends align with the holes in the rack posts.
- Tighten the screws and nuts according to instructions for your rack system.






Figure 2 - 3: Rail ends attach to the outside of each post



1. Place the Vess R2600 subsystem onto the rails.

- At least two persons are required to safely lift the Vess R2600.
- Lift the Vess R2600 subsystem itself. Do not lift the Vess R2600 by its brackets.


Figure 2 - 4: Placing the Vess R2600 subsystem onto the rack rails



1. Secure the Vess R2600 subsystem to the rack.

- The Vess R2600 attaches to the rack posts using the included screws and flange nuts.
- Use the attaching screws and flange nuts that came with the Vess R2600.

Figure 2 - 2: Placing the Vess R2600 subsystem onto the rack rails



INSTALLING PHYSICAL DRIVES

If your Vess R2600 subsystem shipped with the drives installed at the factory, you can skip this section and go to

"Making Serial Cable Connections" on page 29

The Vess R2600 subsystems and JBOD expansion units support:

- SAS and SATA physical drives
- 2.5-inch and 3.5-inch physical drives
- Hard disk drives (HDD) and solid state drives (SSD)

For a list of supported physical drives, download the latest compatibility list from PROMISE support:

<http://www.promise.com/support/>.

NUMBER OF DRIVES REQUIRED


The table below shows the number of drives required for each RAID level.

Level	Number of Drives	Level	Number of Drives
RAID 0	1 or more	RAID 6	4 to 32
RAID 1	2 only	RAID 10	4 or more*
RAID 1E	2 or more	RAID 30	6 or more
RAID 3	3 to 32	RAID 50	6 or more
RAID 5	3 to 32	RAID 60	8 or more
* Must be an even number of drives.			

DRIVE SLOT NUMBERING

You can install any suitable disk drive into any slot in the enclosure. The diagram below shows how Vess R2600's drive slots are numbered. Slot numbering is reflected in the WebPAM PROe and CLU user interfaces.

Figure 2 - 5: Drive slot numbering



Install all of the drive carriers into the Vess R2600 enclosure to ensure proper airflow, even if you do not populate all the carriers with physical drives.

INSTALLING YOUR DRIVES

The drive carrier accommodates 2.5-inch and 3.5-inch drives, with or without a SAS-to-SATA adapter.



Cautions

- Swing open the drive carrier handle before you insert the drive carrier into the enclosure.
- To avoid hand contact with an electrical hazard, remove only one drive carrier a time.




Important

- SATA drives require a SAS-to-SATA adapter, available from PROMISE Technology at <http://www.promise.com>
- SAS drives do not require adapters.

- Press the drive carrier release button. The handle springs open.
- Grasp the handle and gently pull the empty drive carrier out of the enclosure.

Figure 2 - 8: Drive carrier front view



- If you are installing SATA drives, attach a SAS-to-SATA adapter onto the power and data connectors of each drive.
- Carefully lay the drive into the carrier with the power and data connectors facing away from the carrier handle.

- Position the drive in the carrier so the mounting holes line up.
 - 2.5-inch drive mounting screws go through the bottom of the carrier.
 - SAS-to-SATA adapter mounting screws go through the bottom of the carrier.
 - 3.5-inch drive mounting screws go through the sides of the carrier.

Figure 2 - 6: Drive carrier bottom view





Figure 2 - 7: Drive carrier side view



- Insert the screws through the proper holes in the carrier and into the drive or adapter.
 - Use the screws supplied with the Vess R2600 or the SAS-to-SATA adapter.
 - Install four screws per drive.
 - Install two screws per adapter.
 - Snug each screw. Be careful not to over tighten.
- With the drive carrier handle in open position, gently slide the drive carrier into the enclosure.



Important

- Press the release button to push the drive carrier into position.
- Proper drive installation ensures adequate grounding and minimizes vibration. Always attach the drive to the carrier with four screws.

MAKING MANAGEMENT AND DATA CONNECTIONS

Examples of Vess R2600 configurations include:

- "Fibre Channel SAN" on page 18
- "Fibre Channel DAS" on page 20
- "Fibre Channel with JBOD Expansion" on page 21
- "Fibre Channel SAN – No Single Point of Failure" on page 22
- "iSCSI Storage Area Network (SAN)" on page 24
- "iSCSI Direct Attached Storage (DAS)" on page 27
- "iSCSI with JBOD Expansion" on page 28


FIBRE CHANNEL SAN



Important

For a list of supported HBAs, Switches, and SFP transceivers, download the latest compatibility list from PROMISE support:
<http://www.promise.com/support/>.

Figure 2 - 9: FC data and management ports on the RAID controller



A Fibre Channel storage area network (SAN) requires:

- An FC HBA card in each host PC or server
- An SFP transceiver for each connected FC port on the subsystem
- An FC switch
- A network switch

FC SAN DATA PATH

To establish the data path:

1. Connect FC cables between at least one FC data port on each RAID controller and the FC switch.

See "Figure 2 - 10: FC SAN data and management connections".

2. Connect FC cables between the FC switch and the FC HBA cards in both host PCs or servers.

If you have multiple Vess R2600 subsystems, repeat steps 1 and 2 as required.

The Vess R2600 subsystem is shown with SFP transceivers installed.

MANAGEMENT PATH

To establish the management path:

1. Connect Ethernet cables between the Management ports on both RAID controllers and the network switch.


See "Figure 2 - 10: FC SAN data and management connections".

2. Connect Ethernet cables between the network ports on both host PCs or servers and the network switch.

If you have multiple Vess R2600 subsystems, repeat steps 1 and 2 as required.

The Vess R2600 subsystem is shown with SFP transceivers installed.

Figure 2 - 10: FC SAN data and management connections



FIBRE CHANNEL DAS



Important

For a list of supported HBAs, switches, and SFP transceivers, download the latest compatibility list from PROMISE support:
<http://www.promise.com/support/>.

Fibre Channel direct attached storage (DAS) requires:

- An FC HBA card in the host PC or server
- An SFP transceiver for each connected FC port on the subsystem
- A network switch

FC DAS DATA PATH

To establish the data path:

1. Connect an FC cable between a data port on the left RAID controller and the FC HBA card in your host PC or server.

See "Figure 2 - 11: FC DAS data and management connections" on page 20

2. Connect an FC cable between a data port on the right RAID controller and the FC HBA card in your host PC or server.

MANAGEMENT PATH

To establish the management path:


1. Connect Ethernet cables between the Management ports of both RAID controllers and the network switch.

See "Figure 2 - 11: FC DAS data and management connections" on page 20

2. Connect an Ethernet cable between the network port on the host PC or server and the network switch.

The Vess R2600 subsystem is shown with SFP transceivers installed.

Figure 2 - 11: FC DAS data and management connections



FIBRE CHANNEL WITH JBOD EXPANSION

JBOD expansion requires at least one SFF-8088 4X to SFF-8088 4X external SAS cable for each JBOD unit.


To add JBOD units:

1. Connect the SAS expansion port on the left controller of the RAID subsystem to the SAS data IN port on the left I/O module of the first JBOD unit.

See "Figure 2 - 12: Vess R2600fiD with FC JBOD expansion connections"
2. Connect the SAS expansion port on the right controller of the RAID subsystem to the SAS data IN port on the right I/O module of the first JBOD unit.
3. Connect the SAS data OUT port on left I/O module of the first JBOD unit to the SAS data IN port on the left I/O module of the second JBOD unit.
4. Connect the SAS data OUT port on right I/O module of the first JBOD unit to the SAS data IN port on the right I/O module of the second JBOD unit.
5. Connect the remaining JBOD units in the same manner.
 - Keep your data paths organized to ensure redundancy.
 - JBOD expansion supports up to four JBOD units.

The Vess R2600 subsystem is shown with SFP transceivers installed.

Figure 2 - 12: Vess R2600fiD with FC JBOD expansion connections



FIBRE CHANNEL SAN – No SINGLE POINT OF FAILURE

An FC SAN with no single point of failure (NSPF) requires:

- An FC HBA card in each host PC or server
- An SFP transceiver for each connected FC port on the subsystem
- Two SFF-8088 4X to SFF-8088 4X SAS external cables for each JBOD unit
- Two FC switches
- A network switch

FC SAN NSPF DATA PATH

To establish the data path:

1. Connect an FC cable between an FC data port on the left RAID controller and one of the FC switches.
2. Connect an FC cable between an FC data port on the left RAID controller and the other FC switch.
3. Connect an FC cable between an FC data port on the right RAID controller and one of the FC switches.
4. Connect an FC cable between an FC data port on the right RAID controller and the other FC switch.
5. Connect FC cables between one of the FC switches and the FC HBA cards in both of the host PCs or servers.
6. Connect FC cables between the other FC switch and the FC HBA cards in both of the host PCs or servers.

If you have multiple Vess R2600 subsystems, repeat steps 1 through 6 as required.




Important

For a list of supported HBAs, switches, and SFP transceivers, download the latest compatibility list from PROMISE support:
<http://www.promise.com/support/>.

The Vess R2600 subsystem is shown with SFP transceivers installed.

Figure 2 - 13: FC SAN NSPF data connections



FC SAN NSPF MANAGEMENT PATH

To establish the management path:

1. Connect an Ethernet cable between the Management port on each RAID controller and the network switch.

See "Figure 2 - 14: FC SAN NSPF management connections"

2. Connect an Ethernet cable between the network port on each host PC or server and the network switch.

If you have multiple Vess R2600 subsystems, repeat steps 1 and 2 as required.


JBOD EXPANSION

JBOD connections are the same for all FC SAN and DAS configurations.

See"Figure 2 - 12: Vess R2600fiD with FC JBOD expansion connections"

The Vess R2600 subsystem is shown with SFP transceivers installed.

Figure 2 - 14: FC SAN NSPF management connections



iSCSI STORAGE AREA NETWORK (SAN)



Important

For a list of supported HBA NICs and switches, download the latest compatibility list from PROMISE support:
<http://www.promise.com/support/>.

This arrangement requires:


- An iSCSI HBA network interface card (NIC) in the host PC or server
- A GbE network switch
- A standard network switch



Note

Only one iSCSI data cable is required between each RAID controller and the GbE network switch. However, you can attach multiple cables to create redundant data paths or trunking.

Figure 2 - 15: iSCSI data and management ports on the RAID controller



iSCSI SAN DATA PATH

Each Vess R2600 controller has four (4) RJ45 iSCSI data port connectors. See "Figure 2 - 15: iSCSI data and management ports on the RAID controller" on page 24


To establish the data path:

1. Connect Ethernet cables between the iSCSI NIC in both host PCs or servers and the GbE network switch.

See "Figure 2 - 16: iSCSI SAN data connections" on page 25.
2. Connect an Ethernet cable between at least one iSCSI data port on the left RAID controller and the GbE network switch.
3. Connect an Ethernet cable between at least one iSCSI data port on the right RAID controller and the GbE network switch.

If you have multiple Vess R2600 subsystems, host PCs or servers, repeat steps 1 through 3 as required.

Figure 2 - 16: iSCSI SAN data connections



MANAGEMENT PATH

Each Vess R2600 controller has one (1) Ethernet RJ45 management port connector. See "Figure 2 - 15: iSCSI data and management ports on the RAID controller" on page 24.

To establish the management path:


1. Connect Ethernet cables between the network connector on both host PCs or servers and the standard network switch.

See "Figure 2 - 17: iSCSI SAN management connections" on page 26.

2. Connect Ethernet cables between the Management port on both RAID controllers to the standard network switch.

If you have multiple Vess R2600 subsystems, repeat steps 1 and 2.

Figure 2 - 17: iSCSI SAN management connections



iSCSI DIRECT ATTACHED STORAGE (DAS)



Important

For a list of supported HBAs and switches, download the latest compatibility list from PROMISE support:
<http://www.promise.com/support/>.

This arrangement requires:

- An iSCSI HBA network interface card (NIC) in the host PC or server
- A standard network switch

DATA PATH

Each Vess R2600 controller has four (4) RJ45 iSCSI data port connectors. See "Figure 2 - 15: iSCSI data and management ports on the RAID controller" on page 24.

To establish the data path:

1. Connect an Ethernet cable between the iSCSI NIC in the host PC or server and an iSCSI data port on one of the RAID controllers.

See "Figure 2 - 18: iSCSI DAS data and management connections" on page 27.

2. Connect an Ethernet cable between the iSCSI NIC in the host PC or server and an iSCSI data port on the other RAID controller.

MANAGEMENT PATH

Each Vess R2600 controller has one (1) Ethernet RJ-45 management port connector. See "Figure 2 - 15: iSCSI data and management ports on the RAID controller" on page 24.


To establish the management path:

1. Connect an Ethernet cable between the network connector on the host PC or server and the standard network switch.

See "Figure 2 - 18: iSCSI DAS data and management connections" on page 27.

2. Connect Ethernet cables between the standard network switch and the Management ports on both RAID controllers.

Figure 2 - 18: iSCSI DAS data and management connections



iSCSI WITH JBOD EXPANSION

JBOD expansion requires at least one SFF-8088 4X to SFF-8088 4X external SAS cable for each JBOD unit.


To add JBOD units:

1. Connect the SAS expansion port on the left controller of the RAID subsystem to the SAS data IN port on the left I/O module of the first JBOD unit.

See "Figure 2 - 19: Vess R2600iD with iSCSI JBOD expansion connections" on page 28.

2. Connect the SAS expansion port on the right controller of the RAID subsystem to the SAS data IN port on the right I/O module of the first JBOD unit.
3. Connect the SAS data OUT port on left I/O module of the first JBOD unit to the SAS data IN port on the left I/O module of the second JBOD unit.
4. Connect the SAS data OUT port on right I/O module of the first JBOD unit to the SAS data IN port on the right I/O module of the second JBOD unit.
5. Connect the remaining JBOD units in the same manner.
 - Keep your data paths organized to ensure redundancy.
 - JBOD expansion supports up to four JBOD units.

Figure 2 - 19: Vess R2600iD with iSCSI JBOD expansion connections



MAKING SERIAL CABLE CONNECTIONS


Serial communication enables the terminal emulation application on your host PC or server to access the Vess R2600's Command Line Interface (CLI) to set up a network connection. The Vess R2600 package includes one RJ-11-to-DB9 serial data cable for each controller.

To set up a serial cable connection:

1. Attach the RJ-11 end of the serial data cable to the RJ-11 serial connector on one of the RAID controllers.
2. Attach the DB9 end of the serial data cable to a serial port on the host PC or server.

Figure 2 - 20: Serial port connection

Use the RJ-11 serial port on the controller module to establish the serial communication link. The Vess R2600 is shipped with an RJ-11 to DB9 adapter to be used for this purpose.



CONNECTING THE POWER

To power on the Vess R2600 and any connected Vess J2600 system, follow these steps:

1. Connect the power insert for all power supplies on all units.
2. Plug in the power cables on all power cords to a suitable grounded power source.
3. To turn on the power to the Vess J2600 or Vess R2600 units, press the power button on the front of the left handle of the Vess R2600 devices. See "Figure 2 - 23: Left side of the front of the Vess R2600" on page 31 for an illustration of the power button. The Vess R2600 features an automatic JBOD detect and power on sequence mechanism so that all connected Vess J2600 expansion units are powered on in the correct sequence. This feature will first power on the Vess J2600s and then the Vess R2600 units in the correct sequence automatically.

After the powering on the Vess R2600 and Vess J2600 units, check the LEDs to monitor the devices.

LED BEHAVIOR

When the power is switched on, the LEDs on the right handle light up. See "Figure 2 - 22: Right side of the Vess R2600 front LED display" on page 31.

When boot-up is finished and the Vess R2600 is functioning normally:


- When boot-up is finished and the Vess R2600 subsystem is functioning normally:
- Power, Global Enclosure Status, and Global RAID Status LEDs display green continuously.
- Controller Activity LED flashes green when there is controller activity.
- System Heartbeat LED blinks blue (once a second), and repeats the pattern.

See the sections that follow for more details on device LED indicator behavior.

REAR PANEL PSU & COOLING FAN LEDs

The LEDs on the rear panel include LEDs on each cooling fan and each power supply. These LEDs will light green to indicate normal operation. A red LED indicates a problem or unit failure.

Figure 2 - 21: LEDs on Power Supply and Cooling Units



FRONT PANEL LEDs

Figure 2 - 23: Left side of the front of the Vess R2600





Figure 2 - 22: Right side of the Vess R2600 front LED display



LED State	Power	Global Enclosure Status	Global RAID Status	Controller Activity	System Heartbeat
<i>Dark</i>	No power	—		No controller installed	—
<i>Steady Green</i>		All devices normal	All LDs are on line	—	—
<i>Steady Blue</i>	Normal			No activity	
<i>Blinking Blue</i>	—	—	—	—	Normal**
<i>Flashing Blue</i>				I/O Activity	
<i>Flashing Green</i>	—	Locating device	—		—
<i>Amber</i>	—	One or two devices in error	One or more LD is critical; none are offline	—	—
<i>Red</i>	—	Three or more devices in error	One or more LD is offline	—	—

For more information on LEDs, see "Chapter 8: Troubleshooting" on page 201.

* Check the LEDs on the back of the Vess R2600 enclosure.
** Blinks blue once a second

DRIVE STATUS INDICATORS

The Vess R2600 and Vess J2600 spins up the disk drives sequentially to equalize power draw during start-up. After a few moments:

- The Power/Activity LED displays blue when a physical drive is present.
- The Drive Status LED displays green when the physical drive is configured as a member of a disk array or as a spare. When the physical drive is unconfigured, the LED is dark.

See the table below for a more complete description of the Drive carrier LEDs.

Figure 2 - 24: Drive carrier LEDs



Drive Carrier LEDs		
State	Power / Activity	Drive Status
Dark	No drive in carrier	Drive is unconfigured
Steady Blue	Drive is present	—
Flashing Blue	Activity on drive	—
Steady green	—	Drive is configured
Blinking	both LEDs blink	Locator feature
Amber	—	Drive is rebuilding
Red	—	Drive error or failure
* Configured means the physical drive either belongs to an array or it is assigned as a spare drive.		
* Steady means the LED is on.		
* Blinking means a regular on/off pattern.		
* Flashing means intermittent and irregular on/off pattern.		

For more information on LEDs, see "Chapter 8: Troubleshooting" on page 201..

CONTROLLER LEDs

When boot-up is finished and the Vess R2600 subsystem is functioning normally:

- Controller status LEDs display green continuously.
- Ethernet LEDs display green or flash depending on your network connection.
- The FC, iSCSI, SAS, and Expansion LEDs display green or flash during port activity.

Figure 2 - 25: Vess R2600 controller LEDs





Figure 2 - 26: Vess R2600fi controller LEDs



CONTROLLER LED BEHAVIOR

When boot-up is finished and the Vess R2600 subsystem is functioning normally, the Controller status LED displays green continuously; the Management port LEDs display green or flash depending on your network connection; the FC, iSCSI, and SAS Expansion LEDs display green or flash during port activity.

LED	Description
SAS Expansion	Lights green when connected, flashes green when active.
Controller Status	This displays the current operational status of the controller. A steady (unblinking) green light indicates the controller is operational.
Dirty Cache	Blinks amber if cache is dirty, meaning that the controller memory cache contains data, otherwise this is dark.
USB	A steady green light indicates a valid USB connection, this is dark when not connected (no device attached).
iSCSI (2 above each port)	Left LED lights green when connected, flashes green when active, dark if not connected. Right LED indicates connection speed, green is 100 Mbps, amber is 1000 Mbps.
FC ports	See next page

LED	Description		
FC ports (one LED above each port)	Green	Amber	Status
	dark	dark	Wake-up failure (dead board)
	off	on	POST failure (dead board)
	off	blinking slowly	Wake-up failure monitor
	off	blinking rapidly	Failure to POST
	off	flashing	POST in progress
	on	off	Failure while functioning
	on	on	Failure while functioning
	on	2 rapid blinks	Normal, link up, 2 Gb/s
	on	3 rapid blinks	Normal, link up, 4 Gb/s
	on	4 rapid blinks	Normal, link up, 8 Gb/s
	blinking slowly	off	Normal, link down
	blinking slowly	blinking slowly	Offoffline for download
	blinking slowly	blinking rapidly	Restricted offoffline mode (waiting for restart)
	blinking slowly	flashing	Restricted offoffline mode, test active

CHAPTER 3: SETUP

This chapter covers the following topics:

- "Setting-up the Serial Connection" on page 35
- "Vess R2600 Default IP Addresses" on page 36
- "Choosing DHCP or a Static IP Address" on page 36
- "Setting-up Vess R2600 with the CLI" on page 37
- "Setting-up Vess R2600 with the CLU" on page 41
- "Logging into WebPAM PROe" on page 44
- "Creating Disk Arrays and Logical Drives" on page 45
- "Enabling LUN Mapping and Masking" on page 49

SETTING-UP THE SERIAL CONNECTION

The initial connection accesses the serial port using the serial cable connection you made using the RJ-11 to DB9. Use your PC’s terminal emulation program, such as Microsoft HyperTerminal, to access the Command Line Interface (CLI).

You can also use the serial connection to manage the Vess R2600 through the Command Line Utility (CLU).

To make the initial serial connection:

1. Change your terminal emulation application settings to match the following specifications:
 - Bits per second: 115200
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: none
2. Start your PC’s terminal VT100 or ANSI emulation program.
3. Press **Enter** once to launch the CLI.
4. At the Login prompt, type and press **Enter**.
5. At the Password prompt, type **password** and press **Enter**.

The screen displays:

```
login as: administrator
administrator@Vess R2600's password:*****
-----
Promise Vess Command Line Interface (CLI) Utility
Version: 1.00.0000.02 Build Date: Sep 3, 2012
-----
Type help or ? to display all the available commands
Type menu to enter Menu Driven Configuration Utility
-----
administrator@cli>
```

To see the full set of CLI commands, at the administrator@cli> prompt, type `help` and press **Enter**.

To see full information about a specific command, at the administrator@cli> prompt, type `help` followed by the command, then press **Enter**.

```
administrator@cli> help net
```

ABOUT IP ADDRESSES

- Vess R2600 Default IP Addresses (page 36)
- Choosing DHCP or a Static IP Address (page 36)
- Accessing the MAC Address in the CLI (page 36)
- Accessing the MAC Address in the CLU (page 37)

Choosing the appropriate IP addresses is essential to manage your Vess R2600 subsystem over a network. You must change the IP addresses of the subsystems as required for your environment.

Vess R2600 Default IP Addresses

The default **virtual** management port IP addresses are set to:

- IPv4 – 10.0.0.1
- IPv6 – 2001::1

The virtual management port IP address works with either RAID controller, enabling you to access a dual-controller Vess R2600 over your network using a single IP address.

The default **physical** management port IP addresses are set to:

- Controller 1, IPv4 – 10.0.0.2
- Controller 1, IPv6 – 2001::2
- Controller 2, IPv4 – 10.0.0.3
- Controller 2, IPv6 – 2001::3

The physical management port IP address works with only one RAID controller and is used when the controller goes into maintenance mode. For more information, see "Maintenance Mode" on page 211.

CHOOSING DHCP OR A STATIC IP ADDRESS

When you setup your Vess R2600, you have the option of:

- Enabling DHCP and letting your DHCP server assign the IP address to the Vess R2600’s virtual management port.
- Specifying a static IP address for the Vess R2600’s virtual management port.

DHCP is currently supported on IPv4 only. If you use IPv6, you must make your network settings manually.

If you choose to enable DHCP, have your Network Administrator dedicate an IP address for the Vess R2600, linked to the Vess R2600’s MAC address. This action prevents the DHCP server from assigning a new IP address when the Vess R2600 restarts, with the result that users can no longer log in.

ACCESSING THE MAC ADDRESS IN THE CLI

To access the MAC address in the CLI:

At the command prompt, type `net -a list -v` and press **Enter**.

The following information displays:

```
administrator@cli> net -a list -v
-----
ActiveCtrlId: 1   Port: 1
MaxSupportedSpeed: 1000Mbps LinkStatus: Up
ProtocolFamily: IPv4(Enabled)   DHCP: Disabled
IP: 10.0.0.1
IPMask: 0.0.0.0
MAC: 00:01:55:61:18:65
DNS: 0.0.0.0
Gateway: 0.0.0.0

ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: ::
IPMask: ::
MAC: 00:01:55:61:18:65
DNS: ::
Gateway: ::
```

ACCESSING THE MAC ADDRESS IN THE CLU

To access the MAC address in the CLU:

1. At the CLI command prompt, type menu and press **Enter**.

The CLU screen appears.

2. Highlight **Network Management** and press **Enter**.
3. Highlight **IPv4** and press **Enter**.

The following information displays:

```
Active Controller Id: 1      Port Id      : 1
Max Supported Speed : 1000Mbps  Link Status    : Up
Protocol Family   : IPv4
Status            : Enabled
MAC Address       : 00:01:55:61:18:65
DHCP              : Disabled
IP Address        : 10.0.0.1
Subnet Mask       : 0.0.0.0
Gateway IP Address : 0.0.0.0
DNS Server IP Address : 0.0.0.0
```

SETTING-UP VESS R2600 WITH THE CLI

Setting up the Vess R2600 in the CLI includes these actions:

- "Making Subsystem Date and Time Settings" on page 37
- "Virtual Management Port Settings" on page 38
 - "Making Virtual Management Port Settings – Automatically" on page 38
 - "Making Virtual Management Port Settings – Manually under IPv4" on page 38
 - "Making Virtual Management Port Settings – Manually under IPv6" on page 38
- "Maintenance Mode Settings" on page 39
 - "Making Maintenance Mode Settings – Automatically" on page 39
 - "Making Maintenance Mode Settings – Manually under IPv4" on page 40
 - "Making Maintenance Mode Settings – Manually under IPv6" on page 40

MAKING SUBSYSTEM DATE AND TIME SETTINGS

To set the subsystem date and time:

1. Type **date -a mod -d** and the date in yyyy/mm/dd format then press **Enter**.

```
administrator@cli> date -a mod -d 2011/03/25
```

2. Type **date -a mod -t** and the time in hh:mm:ss format, then press **Enter**.

```
administrator@cli> date -a mod -t 14:50:05
```

You can combine date and time settings, such as:

```
administrator@cli> date -a mod -d 2011/03/25 -t 14:50:05
```

VIRTUAL MANAGEMENT PORT SETTINGS

MAKING VIRTUAL MANAGEMENT PORT SETTINGS – AUTOMATICALLY

Automatic settings require a DHCP server on your network. DHCP is currently supported on IPv4 only.

To enable automatic management port settings:

1. At the command prompt, type **net -a mod -f ipv4 -s “dhcp=enable”** and press **Enter**.

```
administrator@cli> net -a mod -f ipv4 -s "dhcp=enable"
```

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. To verify the setting change, at the command prompt, type **net** and press **Enter**. The following information displays:

```
administrator@cli> net
=====
PF   Status   IP           Link
=====
IPv4 Enabled 10.0.0.1    Up
IPv6 Disabled ::         Up
```

In the above example:

- PF refers to IP protocol family, v4 or v6
- Status refers to whether the IP protocol is enabled. IPv4 is enabled by default.
- IP is the virtual management port IP address.
- Link indicates whether there is a working network connection.

By default, IPv4 is enabled and IPv6 is disabled. Currently IPv6 does not suppport DHCP.

MAKING VIRTUAL MANAGEMENT PORT SETTINGS – MANUALLY UNDER IPv4

To make IPv4 settings manually on the management port:

1. At the command prompt, type **net -a mod -f ipv4 -s “** followed by:
 - **primaryip=** and the IP address ,
 - **primaryipmask=** and the subnet mask ,
 - **primarydns=** and the DNS server IP address ,
 - **gateway=** and the Gateway server IP address “ and press **Enter**.

Example:

```
administrator@cli> net -a mod -f ipv4 -s "primaryip=10.0.0.1, primaryipmask=255.255.255.0,primarydns=10.0.0.11,gateway=10.0.0.1"
```

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. To verify the settings, at the command prompt, type **net -a list -v** and press **Enter**.

The following information displays:

```
administrator@cli> net -a list -v
-----
ActiveCtrlId: 1   Port: 1
MaxSupportedSpeed: 1000Mbps LinkStatus: Up
ProtocolFamily: IPv4(Enabled)   DHCP: Disabled
IP: 10.0.0.1
IPMask: 255.255.255.0
MAC: 00:01:55:61:18:65
DNS: 10.0.0.11
Gateway: 10.0.0.1

ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: ::
IPMask: ::
MAC: 00:01:55:61:18:65
DNS: ::
Gateway: ::
```

MAKING VIRTUAL MANAGEMENT PORT SETTINGS – MANUALLY UNDER IPv6

To make IPv6 settings manually on the management port:

1. At the command prompt, type **net -a enable -f ipv6** and press **Enter** to enable IPv6 on the Vess R2600.

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. At the command prompt, type **net -a mod -f ipv6 -s “** followed by:
 - **primaryip=** and the IP address ,
 - **primaryipmask=** and the subnet mask ,
 - **primarydns=** and the DNS server IP address ,
 - **gateway=** and the Gateway server IP address “ and press **Enter**.

Example:

```
administrator@cli> net -a mod -f ipv6 -s
"primaryip=2001:0db8:85a3:0000:0000:8a2e:0370:7334,
primaryipmask=2001:0db8:fedc:ba98:7654:3210:0246:8acf
primarydns=2001:0db8:85a3:0000:0000:8a2e:0370:7001,
gateway=2001:0db8:85a3:0000:0000:8a2e:0370:7002"
```

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

- 3. To verify the settings, at the command prompt, type net -a list -v and press **Enter**.

The following information displays:

```
administrator@cli> net -a list -v
-----
ActiveCtrlId: 1   Port: 1
MaxSupportedSpeed: 1000Mbps LinkStatus: Up
ProtocolFamily: IPv4(Enabled)   DHCP: Disabled
IP: 10.0.0.1
IPMask: 255.255.255.0
MAC: 00:01:55:61:18:65
DNS: 10.0.0.11
Gateway: 10.0.0.1

ProtocolFamily: IPv6(Enabled)   DHCP: Disabled
IP: 2001:0db8:85a3:0000:0000:8a2e:0370:7334
IPMask: 2001:0db8:fedc:ba98:7654:3210:0246:8acf
MAC: 00:01:55:61:18:65
DNS: 2001:0db8:85a3:0000:0000:8a2e:0370:7001
Gateway: 2001:0db8:85a3:0000:0000:8a2e:0370:7002
```

MAINTENANCE MODE SETTINGS

You also have the option to make maintenance mode settings at a later time in WebPRM PROe.

MAKING MAINTENANCE MODE SETTINGS – AUTOMATICALLY

Automatic settings require a DHCP server on your network. DHCP is currently supported on IPv4 only.

You make maintenance mode settings for one controller at a time.

To enable automatic maintenance mode settings:

- 1. At the command prompt, type net -a mod -m -c 1 -f ipv4 -s "dhcp=enable" and press **Enter**.
administrator@cli> net -a mod -m -c 1 -f ipv4 -s "dhcp=enable"

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

- 2. To verify the settings changes, at the command prompt, type net -a list -m and press **Enter**.

The following information displays:

```
administrator@cli> net -a list -m
-----
CtrlId: 1   Port: 1
ProtocolFamily: IPv4(Enabled)   DHCP: Enabled
IP: 10.0.0.2
IPMask: 255.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 1   Port: 1
ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: 2001::2
IPMask: ffff::
MAC: 00:01:55:30:65:E9
DNS: ::
Gateway: ::

CtrlId: 2   Port: 1
ProtocolFamily: IPv4(Enabled)   DHCP: Disabled
IP: 10.0.0.3
IPMask: 255.0.0.0
MAC: 00:01:55:30:65:D7
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 2   Port: 1
ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: 2001::3
IPMask: ffff::
MAC: 00:01:55:30:65:D7
DNS: ::
Gateway: ::
```

- 3. Repeat steps 1 and 2 above but change -c 1 (controller 1) to -c 2 (contoller 2).

MAKING MAINTENANCE MODE SETTINGS – MANUALLY UNDER IPv4

You make these settings for one controller at a time.

To make maintenance mode settings:

1. At the command prompt, type `net -a mod -m -c 1 -s "` followed by:
 - **primaryip=** and the IP address ,
 - **primaryipmask=** and the subnet mask ,
 - **primarydns=** and the DNS server IP address ,
 - **gateway=** and the Gateway server IP address
“ and press **Enter**.

Example:

```
administrator@cli> net -a mod -m -c 1 "primaryip=10.0.0.101, primaryipmask=255.255.255.0,primarydns=10.0.0.11,gateway=10.0.0.1"
```

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. To verify the settings changes, at the command prompt, type `net -a list -m` and press **Enter**. The following information displays:

```
administrator@cli> net -a list -m
-----
CtrlId: 1          Port: 1
ProtocolFamily: IPv4(Enabled) DHCP: Disabled
IP: 10.0.0.2
IPMask: 255.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 1          Port: 1
ProtocolFamily: IPv6(Disabled) DHCP: Disabled
IP: 2001::2
IPMask: ffff::
MAC: 00:01:55:30:65:E9
DNS: ::
Gateway: ::

CtrlId: 2          Port: 1
ProtocolFamily: IPv4(Enabled) DHCP: Disabled
IP: 10.0.0.3
IPMask: 0.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 2          Port: 1
ProtocolFamily: IPv6(Disabled) DHCP: Disabled
IP: 2001::3
IPMask: ffff::
MAC: 00:01:55:30:65:D7
DNS: ::
Gateway: ::
```

3. Repeat steps 1 and 2 above but change -c 1 (controller 1) to -c 2 (controller 2).

MAKING MAINTENANCE MODE SETTINGS – MANUALLY UNDER IPv6

You make these settings for one controller at a time.

To make maintenance mode settings:

1. At the command prompt, type `net -a enable -f ipv6 -m -c 1` and press **Enter** to enable IPv6.

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. At the command prompt, type `net -a mod -m -c 1 -s "` followed by:

- **primaryip=** and the IP address ,
- **primaryipmask=** and the subnet mask ,
- **primarydns=** and the DNS server IP address ,
- **gateway=** and the Gateway server IP address
“ and press **Enter**.

Example:

```
administrator@cli> iscsi -a mod -t portal -s "primaryip=2001:0db8:85a3:0000:0000:8a2e:0370:7336, primaryipmask=2001:0db8:fedc:ba98:7654:3210:0246:8acf, primarydns=2001:0db8:85a3:0000:0000:8a2e:0370:7001, gateway=2001:0db8:85a3:0000:0000:8a2e:0370:7002"
```

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

3. To verify the settings, at the command prompt, type net -a list -m and press **Enter**.

The following information displays:

```
administrator@cli> net -a list -m
-----
CtrlId: 1          Port: 1
ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
IP: 10.0.0.2
IPMask: 255.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 1          Port: 1
ProtocolFamily: IPv6(Enabled)    DHCP: Disabled
IP: 2001:0db8:85a3:0000:0000:8a2e:0370:7336
IPMask: 001:0db8:fedc:ba98:7654:3210:0246:8acf
MAC: 00:01:55:30:65:E9
DNS: 2001:0db8:85a3:0000:0000:8a2e:0370:7001
Gateway: 2001:0db8:85a3:0000:0000:8a2e:0370:7002

CtrlId: 2          Port: 1
ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
IP: 10.0.0.3
IPMask: 0.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 2          Port: 1
ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: 2001::3
IPMask: ffff::
MAC: 00:01:55:30:65:D7
DNS: ::
Gateway: ::
```

4. Repeat steps 1, 2, and 3 above but change **-c 1** (controller 1) to **-c 2** (controller 2).

This completes management port and maintenance mode setup.

SETTING-UP VESS R2600 WITH THE CLU

Setting up the Vess R2600 in the CLU includes these actions:

- "Accessing the CLU Quick Setup Menu" on page 42
- "Making Subsystem Date and Time Settings" on page 42
- "Virtual Management Port Settings" on page 42
 - "Making Virtual Management Port Settings – Automatically" on page 42
 - "Viewing Virtual Management Port Settings" on page 43
 - "Making Virtual Management Port Settings – Manually under IPv4" on page 43
 - "Making Virtual Management Port Settings – Manually under IPv6" on page 43
- "Maintenance Mode Settings" on page 43
 - "Making Maintenance Mode Settings – Automatically" on page 43
 - "Making Maintenance Mode Settings – Manually under IPv4" on page 43
 - "Making Maintenance Mode Settings – Manually under IPv6" on page 44
- "Exiting the CLU" on page 44

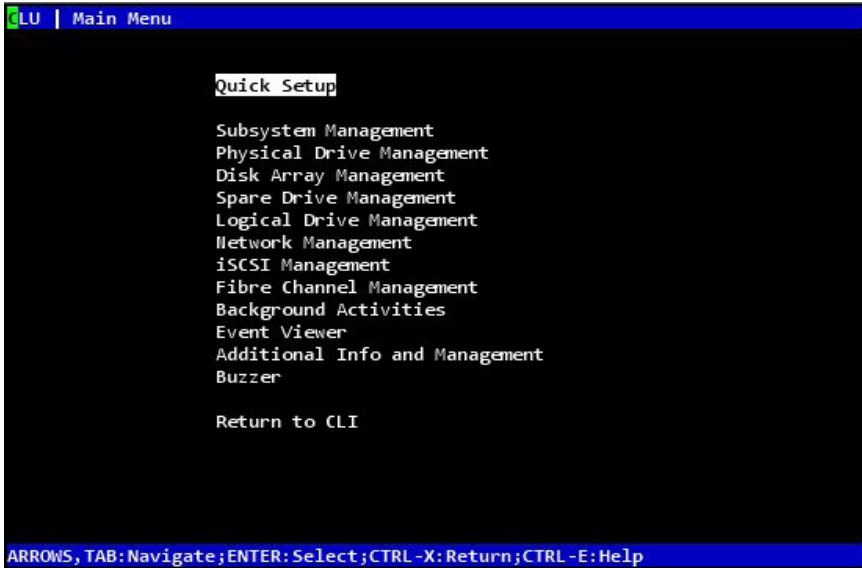
ACCESSING THE CLU QUICK SETUP MENU

To access the Quick Setup menu in the command line utility:

1. At the administrator@cli> prompt, type **menu** and press **Enter**.

The CLU main menu appears. See Figure 3.

Figure 3 - 1: CLU main menu



2. Highlight Quick Setup and press **Enter**.

The first Quick Setup screen enables you to make Date and Time settings.

MAKING SUBSYSTEM DATE AND TIME SETTINGS

To set the subsystem date and time:

1. Press the arrow keys to highlight System Date.
2. Press the backspace key to erase the current date.
3. Type the new date.
4. Follow the same procedure to set the System Time.
5. Press Control-A to save these settings and move to the Management Port settings screen.

VIRTUAL MANAGEMENT PORT SETTINGS

MAKING VIRTUAL MANAGEMENT PORT SETTINGS – AUTOMATICALLY

Automatic settings require a DHCP server on your network. DHCP is currently supported on IPv4 only. Under Quick Setup, the Management Port IPv4 settings screen follows the System Date and Time settings screen.

To enable automatic management port settings:

1. Press the arrow keys to highlight DHCP.
2. Press the spacebar to toggle to Enable.
3. Press Control-A to save these settings and move to the Management Port IPv6 settings screen.

VIEWING *VIRTUAL MANAGEMENT PORT SETTINGS*

To view the current IP address and network settings when using DHCP:

1. Press the arrow keys to highlight DHCP.
2. Press the spacebar to toggle to Disable.

The following information displays:

```
IP Address      : 10.0.0.1
Subnet Mask     : 255.0.0.0
Gateway IP Address : 0.0.0.0
DNS Server IP Address : 0.0.0.0
```

3. Press the spacebar to toggle DHCP back to Enable.
4. Press Control-A to save these settings and move to the Management Port IPv6 settings screen.

MAKING VIRTUAL MANAGEMENT PORT SETTINGS – MANUALLY UNDER IPv4

To make IPv4 settings manually on the management port:

1. Press the arrow keys to highlight IP Address.
2. Press the backspace key to erase the current IP address.
3. Type the new Management Port IP address.
4. Follow the same procedure to specify the Subnet Mask, Gateway IP Address and DNS Server IP Address.

If you do not have a DNS server, skip the DNS Server IP address.

5. Press Control-A to save your settings and move to the Management Port IPv6 settings screen.

MAKING VIRTUAL MANAGEMENT PORT SETTINGS – MANUALLY UNDER IPv6

To make IPv6 settings manually on the management port:

1. Press the arrow keys to highlight IP Address.
2. Press the backspace key to erase the current IP address.
3. Type the new Management Port IP address.
4. Follow the same procedure to specify the Subnet Mask, Gateway IP Address and DNS Server IP Address.

If you do not have a DNS server, skip the DNS Server IP address.

5. Press Control-A to save your settings and move to the Maintenance Mode screens.

MAINTENANCE MODE SETTINGS

For more information, see "Maintenance Mode" on page 211.

You have the option to make maintenance mode settings at a later time in WebPAM PROe.

Under Quick Setup, maintenance mode settings are made in the following sequence:

1. Controller 1, IPv4
2. Controller 1, IPv6
3. Controller 2, IPv4
4. Controller 2, IPv6

MAKING MAINTENANCE MODE SETTINGS – AUTOMATICALLY

Automatic settings require a DHCP server on your network. DHCP is currently supported on IPv4 only.

To enable automatic maintenance mode settings:

1. From the CLU Main Menu, highlight Network Management and press **Enter**.
2. Highlight Maintenance Mode Network Configuration and press **Enter**.
3. Highlight the controller you want and press **Enter**.
4. Highlight DHCP and press the spacebar to toggle to Enabled.
5. Press Control-A to save your settings and move to the Maintenance Mode IPv6 settings screen.

MAKING MAINTENANCE MODE SETTINGS – MANUALLY UNDER IPv4

To make maintenance mode IPv4 manual settings:

1. From the CLU Main Menu, highlight Network Management and press **Enter**.
2. Highlight Maintenance Mode Network Configuration and press **Enter**.
3. Highlight the controller you want and press **Enter**.
4. Highlight DHCP and press the spacebar to toggle to Disabled.
5. Highlight each of the following and press the backspace key to erase the current value, then type the new value.
 - IP address
 - Subnet Mask
 - Default Gateway IP address
 - DNS Server IP address

6. Press Control-A to save your settings and move to the Maintenance Mode IPv6 settings screen.

MAKING MAINTENANCE MODE SETTINGS – MANUALLY UNDER IPv6

To make maintenance mode IPv6 manual settings:

1. From the CLU Main Menu, highlight Network Management and press **Enter**.
2. Highlight Maintenance Mode Network Configuration and press **Enter**.
3. Highlight the controller you want and press **Enter**.
4. Highlight DHCP and press the spacebar to toggle to Disabled.
5. Highlight each of the following and press the backspace key to erase the current value, then type the new value.
 - IP address
 - Subnet Mask
 - Default Gateway IP address
 - DNS Server IP address
6. Press Control-A to save your settings and:
 - If you made settings for Controller 1, move to the Maintenance Mode settings for Controller 2.
 - If you made settings for Controller 2, move to the RAID Configuration menu.



Note

If you want to configure your RAID system now, using the CLU, see "Managing Disk Arrays (CLU)" on page 132 for information about your choices.

EXITING THE CLU

To exit the CLU from the Quick Setup RAID Configuration menu:

1. Highlight Skip the Step and Finish and press **Enter**.
2. Highlight **Return to CLI** and press **Enter**.

This completes management port and maintenance mode setup. Go to Logging into WebPAM PROe.

SETTING UP VESS R2600 WITH WEBPAM PROE

LOGGING INTO WEBPAM PROE

1. Launch your browser.
2. In the browser address field, type in the virtual management port IP address of the Vess R2600 subsystem.

Use the virtual management port IP address you set in the CLI (Setting-up Vess R2600 with the CLI) or CLU (Setting-up Vess R2600 with the CLU). Example:

- WebPAM PROe uses a secure HTTP connection.https://
- **Enter** the IP address of the Vess R2600 10.0.0.1

Together, your entry looks like this: **https://10.0.0.1**

3. When the log-in screen appears:
 - Type **administrator** in the User Name field.
 - Type **password** in the Password field.

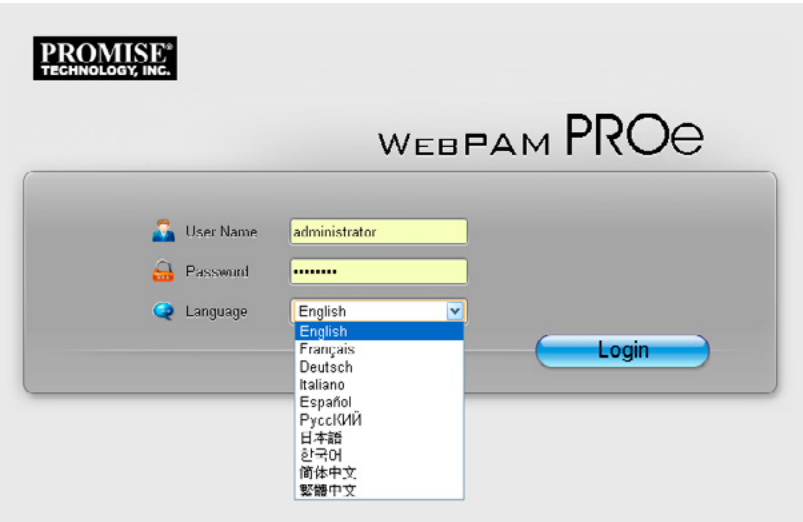
The User Name and Password are case sensitive.

4. Optional. Choose a display language from the drop-down menu.

WebPAM PROe displays in English, German, French, Italian, Spanish, Russian, Japanese, Traditional Chinese, Simplified Chinese, and Korean.

5. Click the **Login** button.

Figure 3 - 2: WebPAM PROe log-in screen with display language options





Important

PROMISE recommends that you change the Administrator’s default password immediately after setup is completed. See "Changing User Passwords" on page 70.

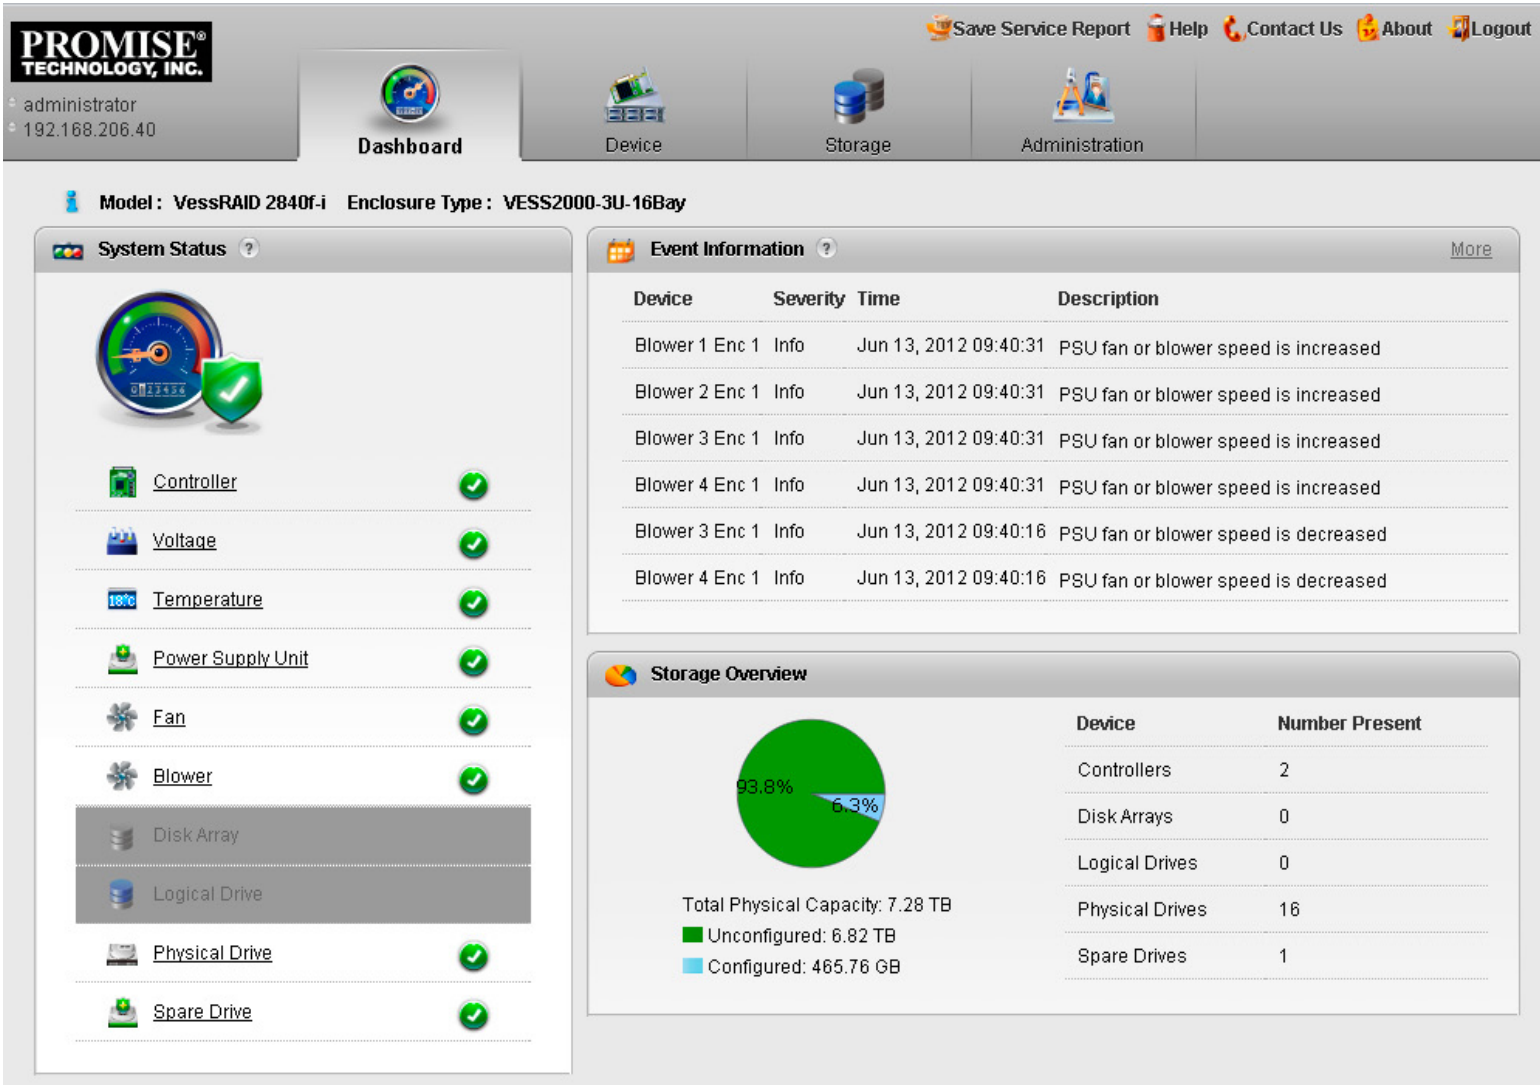


Note

Make a Bookmark (Firefox) or set a Favorite (Internet Explorer) of the Login Screen so you can access it easily next time.

After log-in, the WebPAM PROe opens with the Dashboard tab.

Figure 3 - 3: WebPAM PROe Dashboard tab



CREATING DISK ARRAYS AND LOGICAL DRIVES

On a newly activated RAID system, there are no disk arrays or logical drives. The term “disk array” includes arrays composed of hard disk drives or solid state drives.

To create your disk arrays and logical drives:

1. Click the Storage tab, then click the Wizard option.

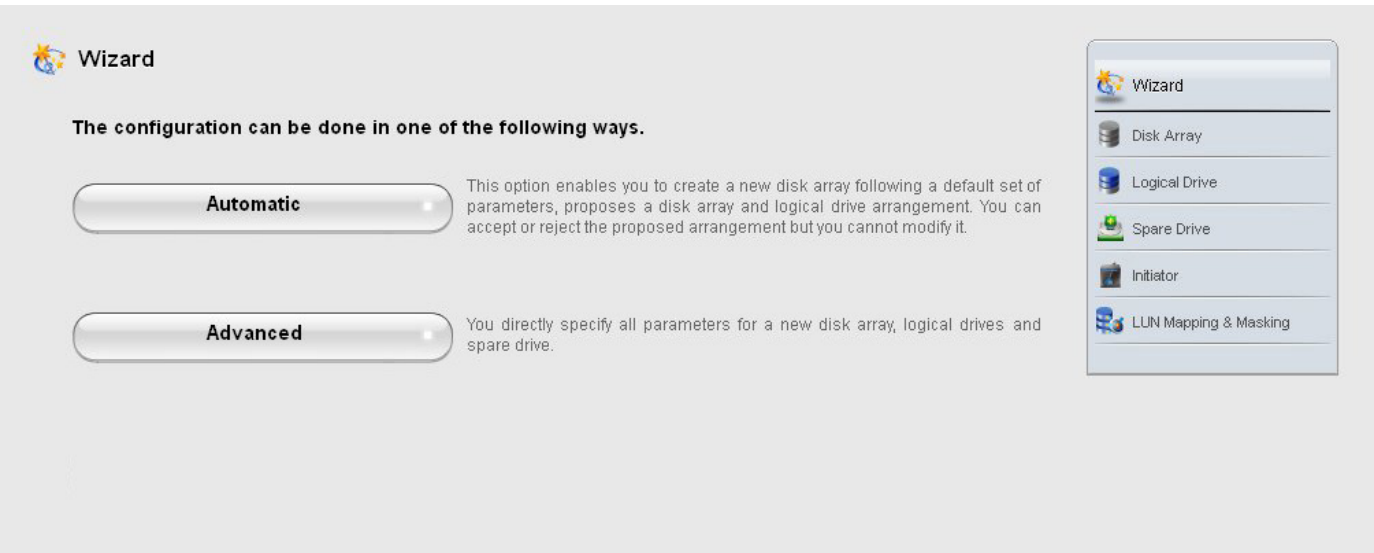
Or, click Disk Array under System Status.

The Wizard screen appears with three creation alternatives:

- Automatic
- Advanced

2. Click one of these buttons to continue.

Figure 3 - 4: The Disk Configuration Wizard main menu



WIZARD - ADVANCED CONFIGURATION



Note

For an explanation of the parameters under the Advanced option, see "Chapter 7: Technology Background" on page 183.

When you choose the **Advanced** option, the **Create Disk Array** menu appears.

STEP 1 – DISK ARRAY CREATION

1. Enter your information and choose your options.
 - Enter a disk array alias in the field provided.
 - Check the box to enable Media Patrol
 - Check the box to enable Predictive Data Migration (PDM)
 - Check the box to enable Power Management
 - Choose a media type – Hard disk drive (HDD) or solid state drive (SSD)

2. Click the enclosure graphic to view information about physical drives.

Look for drives with a green LED dark, a blue LED lit, and no crosshatching over the carrier.

3. Click a physical drive to select it for your array.

The physical drive’s ID number is added to the Selected list.

4. Click the **Next** button to continue.

The **Create Logical Drive** screen appears.

Figure 3 - 6: Create Disk Array menu

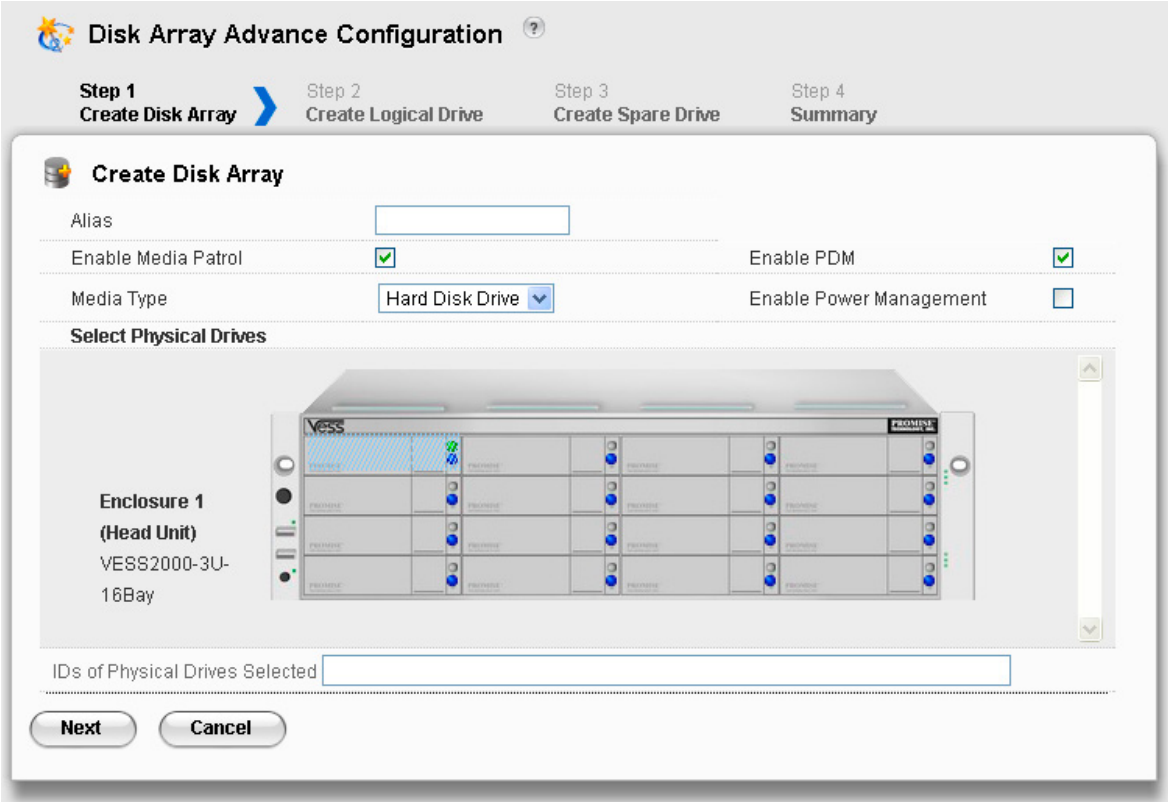
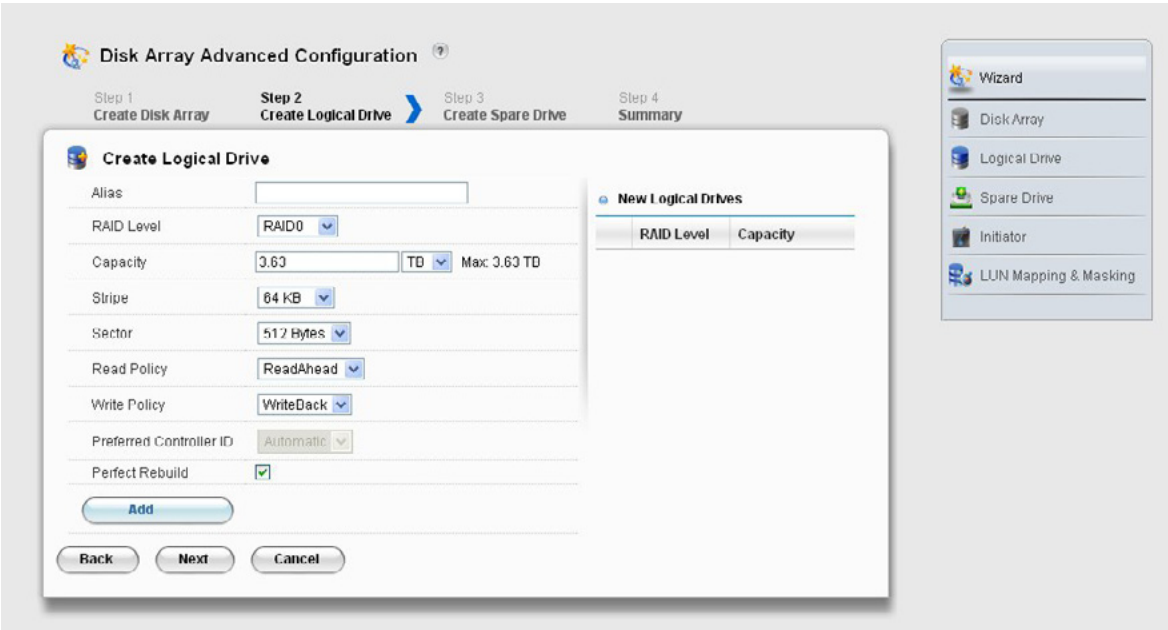


Figure 3 - 7: Create Logical Drive menu



STEP 2 – LOGICAL DRIVE CREATION

1. Enter your information and choose your options.
- Enter a logical drive alias in the field provided.
 - Choose a RAID level from the drop-down menu.
The choice of RAID levels depends on the number of physical drives in your array.
 - Note the **Max:** capacity value. Then enter a capacity value the field provided and choose a unit of measure from the drop-down menu.
 - Choose a stripe size from the drop-down menu.
The choices are 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB.
 - Choose a sector size from the drop-down menu.
The choices are 512 B, 1 KB, 2 KB, and 4 KB.
 - Choose the Read Cache Policy from the drop-down menu
The choices are Read Cache, Read Ahead (cache), and None.
 - Choose the Write Cache Policy from the drop-down menu - The choices are WriteThru (write through) and WriteBack. Write back requires a Read Cache or Read Ahead Read Cache Policy.
 - Uncheck the Perfect Rebuild check box if do not need perfect rebuild for this LD.
 - Click the **Add** button to continue.

The logical drive you just created appears in the **New Logical Drives** list.

2. Click the **Next** button to continue.
- The **Create Spare Drive** screen appears.

STEP 3 – SPARE DRIVE CREATION

Creating a spare drive is optional but highly recommended.

1. Enter your information and choose your options.
- Check the **Revertible** box if you want this spare drive to be revertible.
 - Choose the option for the type spare drive you want.
Global – Replaces a failed drive in any disk ar ray.
Dedicated – Replaces the failed drive only in the assigned disk array.

2. Click the enclosure graphic to view information about physical drives.
3. Click a physical drive to select it for your spare drive.

The physical drive’s ID number is added to the Selected list.

4. Click the **Next** button to continue.
- The **Summary** screen appears.

Figure 3 - 8: Create Spare Drive menu

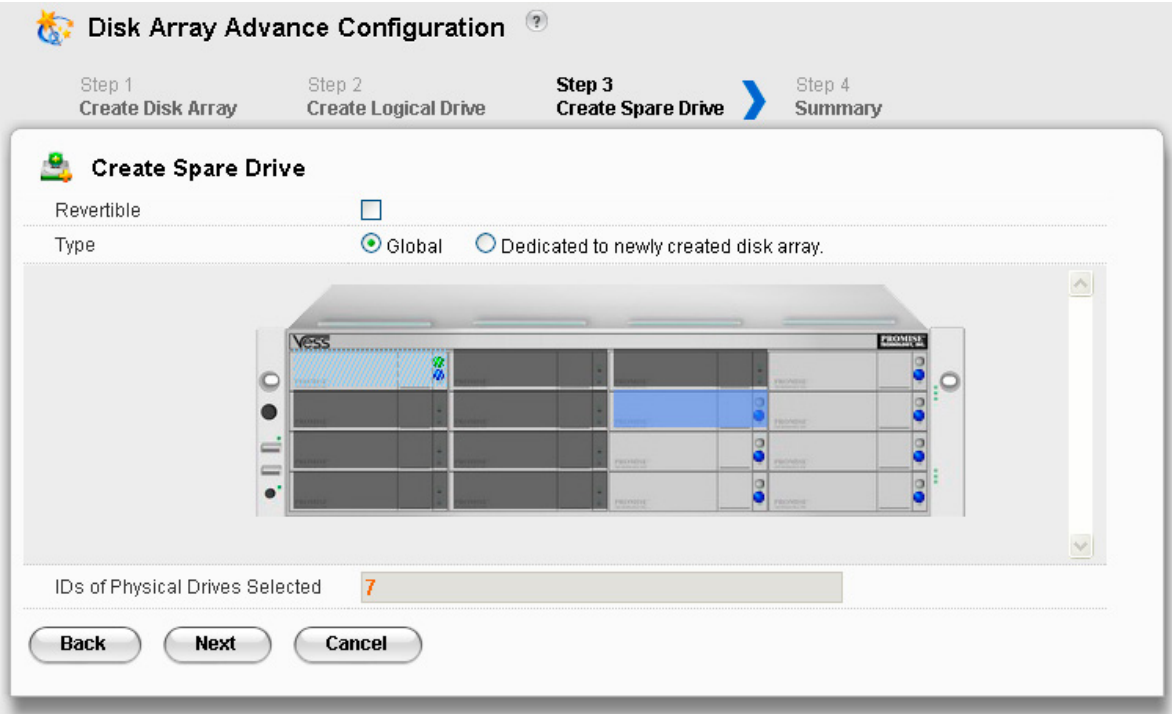
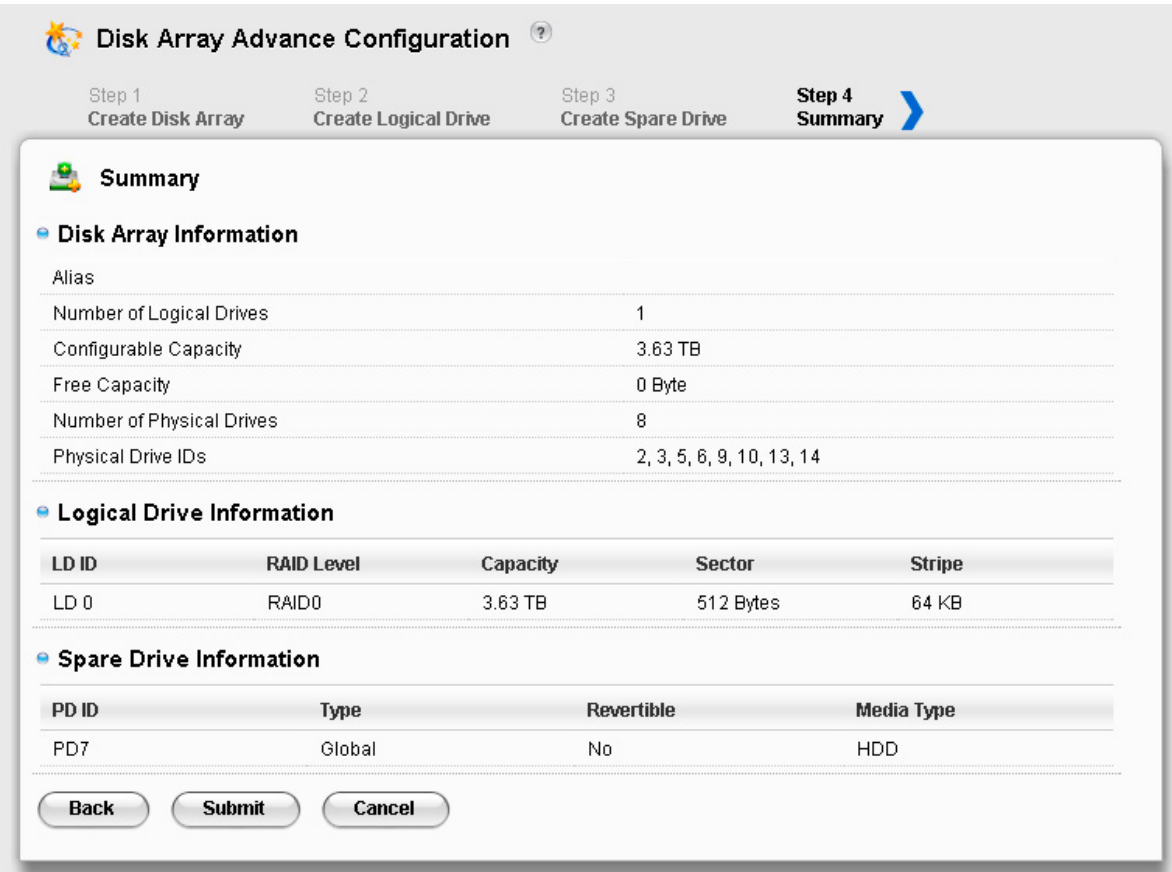


Figure 3 - 9: Disk Array Configuration Summary



STEP 4 – SUMMARY

The Summary screen lists the disk arrays, logical drives, and spare drives that you specified.

If you accept these parameters, click the **Submit** button.

If you do NOT accept these parameters, review and modify your selections in the previous steps.

ENABLING LUN MAPPING AND MASKING

These features are optional for each logical drive. The Enable LUN Mapping dialog box appears after you create a logical drive.

To enable LUN Mapping:

1. Click the **OK** button in the **Enable LUN Mapping** dialog box.

The **LUN Mapping & Masking** screen appears.

2. Check the **Enable LUN Masking** box to enable LUN Masking.
3. Click the **LUN Mapping** button to continue.

The initiator list screen displays.

4. Choose the initiators you want to use from the drop-down menu and click the **Next** button.

The screen displays a list of initiators and a list of logical drives.

5. Click and drag a logical drive from the logical drives list to the initiators list.
6. Click the **Next** button when you are done.

The screen displays a list of initiator IDs and corresponding LUN maps that you specified.

7. Click the **Submit** button to create the LUN map.

The screen displays a list of initiator IDs and corresponding LUN maps.

You can also set LUN mapping and masking at a later time. Click the **Administration** tab, then click the **LUN Mapping & Masking** option.

LOGGING OUT OF WEBPAM PROE

There are two ways to log out of WebPAM PROe:

- Close your browser window
- Click **Logout** on the WebPAM PROe banner

Clicking **Logout** brings you back to the Login Screen. See WebPAM PROe log-in screen with display language options.

After logging out, you must enter your user name and password in order to log in again.

USING WEBPAM PROE OVER THE INTERNET

The above instructions cover connections between Vess R2600 and your company network. It is also possible to connect to a Vess R2600 from the Internet.

Your MIS Administrator can tell you how to access your network from outside the firewall. Once you are logged onto the network, you can access the Vess R2600 using its IP address. See "Logging into WebPAM PROe" on page 44.

CHAPTER 4: MANAGE WITH WEBPAM PROe

This chapter contains the following topics:

- "Logging into WebPAM PROe" on page 50
- "Choosing the Display Language" on page 51
- "Perusing the Interface" on page 51
- "Logging out of WebPAM PROe" on page 52
- "Viewing the Storage Network" on page 53
- "Managing Subsystems" on page 53
- "Managing RAID Controllers" on page 58
- "Managing Enclosures" on page 63
- "Managing UPS Units" on page 67
- "Managing Network Connections" on page 68
- "Managing UPS Units" on page 67
- "Managing Background Activities" on page 72
- "Managing Storage Services" on page 78
- "Monitoring Performance" on page 85
- "Managing Physical Drives" on page 87
- "Managing Disk Arrays" on page 91
- "Managing Logical Drives" on page 96
- "Managing Spare Drives" on page 101
- "Managing Initiators" on page 104
- "Managing LUNs" on page 105
- "Managing Fibre Channel Connections" on page 107
- "Managing iSCSI Connections" on page 109

LOGGING INTO WEBPAM PROe

1. Launch your browser.
2. In the browser address field, type in the virtual management port IP address of the Vess R2600 subsystem.

Use the IP address you set in the CLI (page 37) or CLU (page 41).

Example:

- WebPAM PROe uses a secure HTTP connection.https://
- **Enter** the IP address of the Vess R2600

For example, if your Vess R2600 has an IP address: 10.0.0.1 your entry looks like this: **https://10.0.0.1**

3. When the login screen appears:
 - Type **administrator** in the User Name field.
 - Type **password** in the Password field.
 - Click the **Login** button.

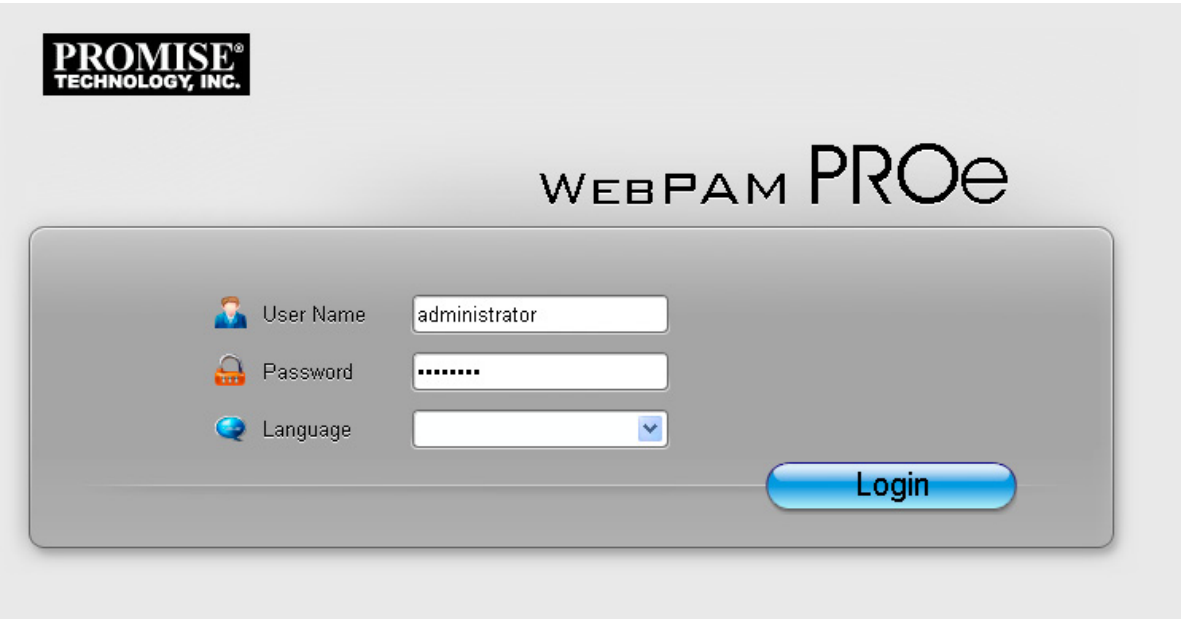
The User Name and Password are case sensitive.

4. Optional. Choose a display language from the drop-down menu.

WebPAM PROe displays in English, German, French, Italian, Spanish, Russian, Japanese, Traditional Chinese, Simplified Chinese, and Korean.

5. Click the Login button.

Figure 4 - 1: WebPAM PROe log-in screen



CHOOSING THE DISPLAY LANGUAGE

WebPAM PROe displays in multiple languages. You choose the display language when you log in.

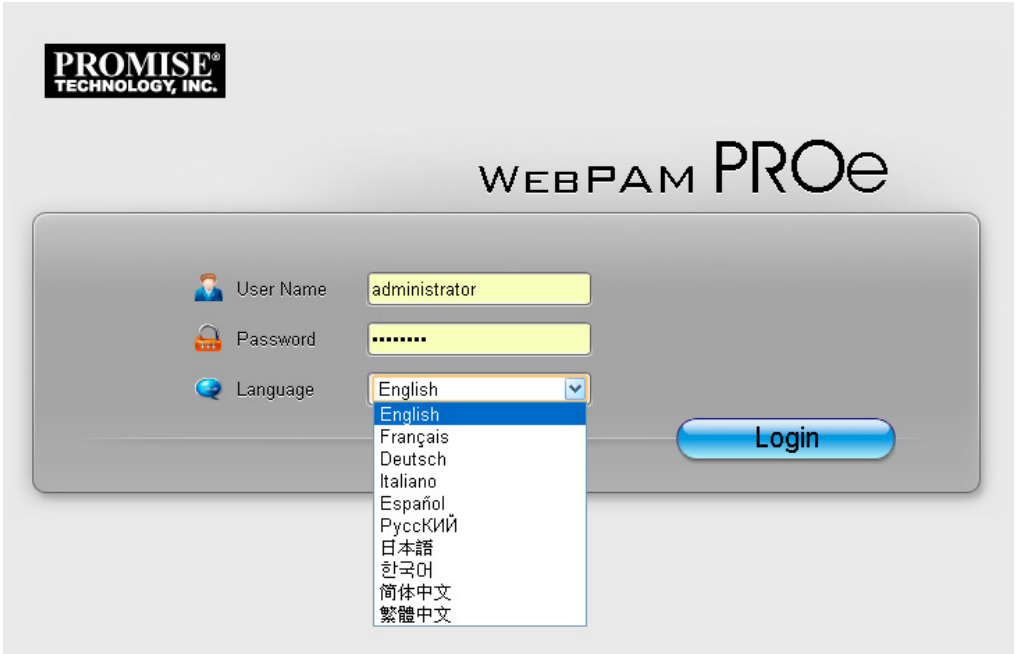
If you are already logged in and you want to change the display language:

1. Click **Logout** at the top right corner of the screen.

The Login screen appears.

2. Click the Language drop-down menu and highlight the language you prefer.

Figure 4 - 2: Login language selection menu



3. Reenter your user name and password.

4. Click the **Login** button.

WebPAM PROe opens in the language you chose.

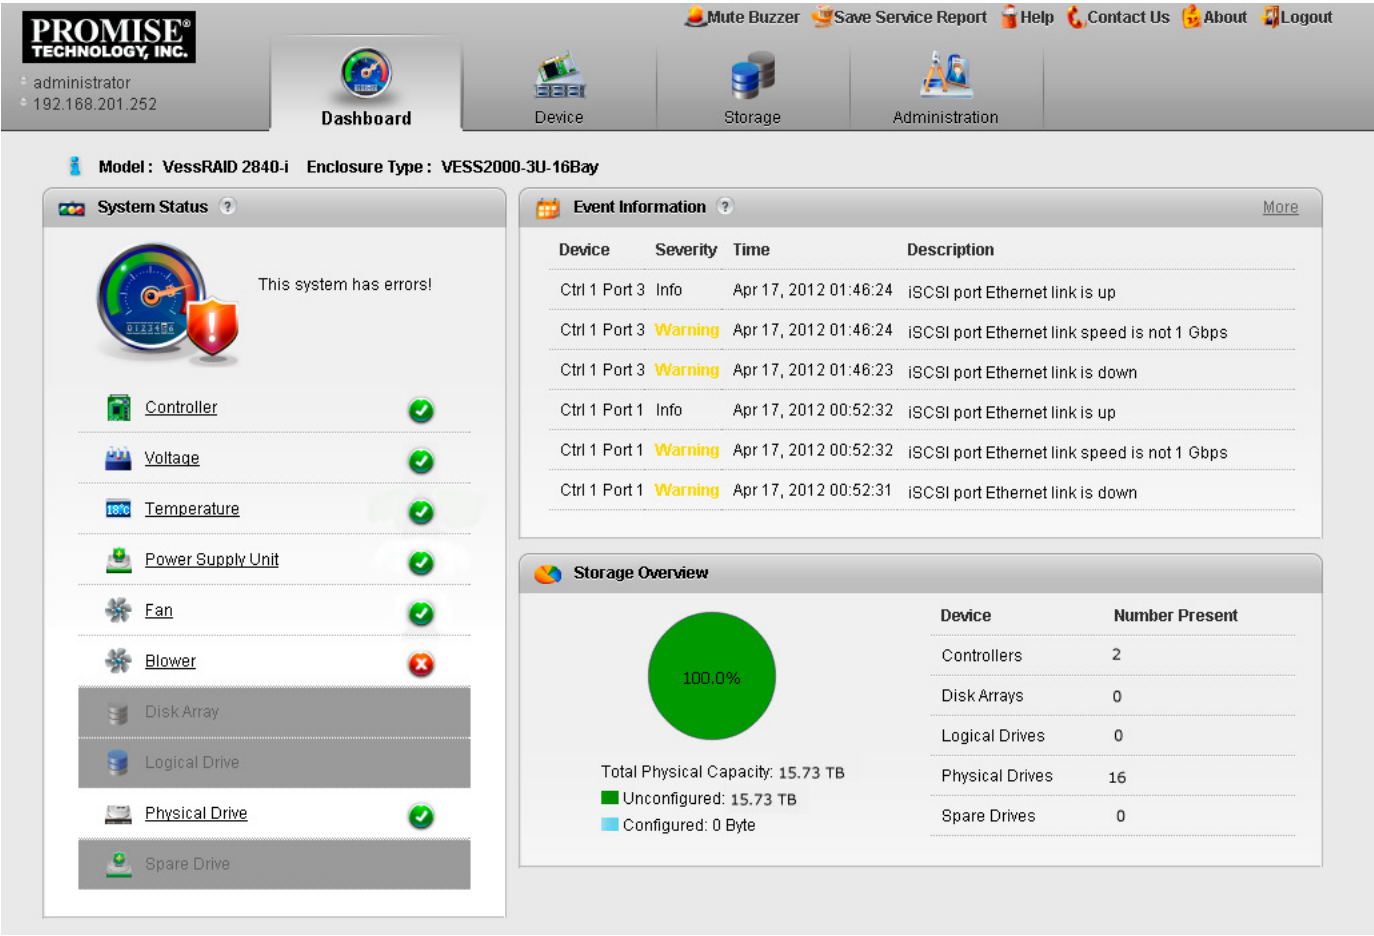
PERUSING THE INTERFACE

The WebPAM PROe interface consists of a header and four tabs, each with specific functions.

- Header
 - Top left corner of the window:
 - Name of logged-in user
 - IP address – Virtual IP address of the RAID subsystem
 - Top right corner of the window
 - Save Service Report – Saves a detailed report to your Host PC
 - Help – Accesses the Help Welcome screen
 - Contact Us – Technical support contact information
 - About – Information about WebPAM PROe
 - Logout – Exits WebPAM PROe
- Discovery tab
 - Displays other PROMISE RAID systems on your network
 - Enables direct login to other PROMISE RAID systems
- Dashboard tab
 - RAID subsystem model and type of enclosure
 - System status
 - Event information – Most recent NVRAM events
 - Storage overview – Capacities, number of devices
- Device tab
 - Enclosure front and back views
 - Topology
 - Enclosure component list and settings
 - Physical drive management
 - UPS (unlimited power supply) management
 - Fibre Channel or iSCSI management
- Storage tab
 - Wizard – Automatic or Advanced configuration
 - Disk array management
 - Logical drive management
 - Initiator management
 - LUN mapping and masking
- Administration tab

- Subsystem settings, clearing statistics, NTP, and controller lock
- User management, including LDAP and role mapping
- Software services
- Runtime and NVRAM event logs
- Background activity, settings and schedules
- Firmware updates
- Image version
- Performance monitor
- PSU wattage monitor
- Restore factory default settings
- Import/Export user database and configuration script
- Network management

Figure 4 - 3: Web PAM PROe Main menu/Dashboard



LOGGING OUT OF WEBPAM PROe

There are two ways to log out of WebPAM PROe:

- Close your browser window
- Click **Logout** on the WebPAM PROe banner

Clicking Logout brings you back to the Login Screen. See WebPAM PROe log-in screen.

After logging out, you must enter your user name and password in order to log in again.

VIEWING THE STORAGE NETWORK

To view the other subsystems on your Storage Network, click the **Discovery** tab at the left edge of the WebPAM PROe window.

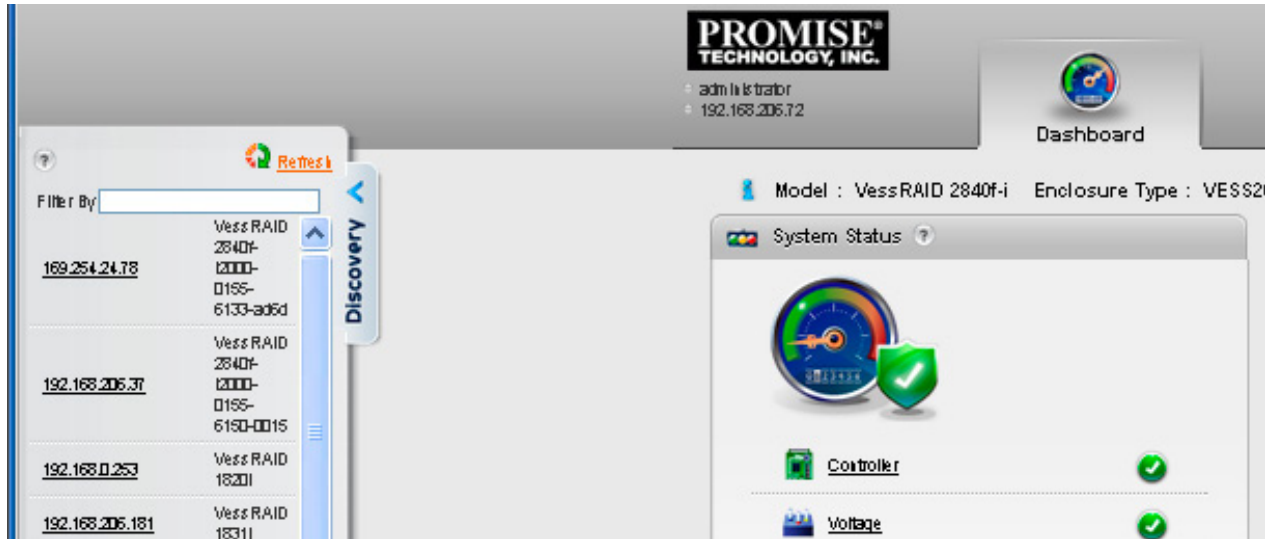
LOGGING ONTO A SUBSYSTEM

To log onto a subsystem in the list, double-click the subsystem.



Caution
The new subsystem displays in the same browser tab. Click your browser’s back button to return to the original subsystem.

Figure 4 - 4: Discovery tab in Main menu



FILTERING THE SUBSYSTEM LIST

To filter the list, so it shows only specific subsystems, enter a characteristic into the Filter By field and press **Enter**. To filter by IP address, enter the IP routing prefix for the range you want to display. For example, typing "10.0" in the entry field reveals all subsystems with IP address beginning with the "10.0" prefix.

REFRESHING THE LIST

To refresh the list, click the **Refresh** link.

MANAGING SUBSYSTEMS

Subsystem management includes:

- "Viewing Subsystem Information" on page 54
- "Making Subsystem Settings" on page 54
- "Locking or Unlocking the Subsystem" on page 54
- "Restoring Factory Default Settings" on page 55
- "Clearing Statistics" on page 55
- "Saving a Service Report" on page 56
- "Importing a Configuration Script" on page 57
- "Exporting a Configuration Script" on page 57
- "Restarting the Subsystem" on page 57
- "Shutting Down the Subsystem" on page 58
- "Restarting the Subsystem after a Shutdown" on page 58

VIEWING SUBSYSTEM INFORMATION

To view subsystem information, click the **Administration** tab.

The list of subsystems and host controllers displays.

Subsystem information includes:

- Alias, if assigned
- Vendor
- Model
- WWN – World Wide Name
- Serial number
- Part number
- Revision number
- Number of JBOD expansion units connected
- Maximum number of JBOD expansion units supported
- Number of controllers present
- Maximum number of controllers supported
- Redundancy status
- Redundancy type
- System date and time

MAKING SUBSYSTEM SETTINGS

To make subsystem settings:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Settings** button.
4. Make changes as required:
 - **Enter** an alias or change the existing alias in the field provided.
 - Choose a redundancy type from the drop-down menu.
The choices are **Active-Active** and **Active-Standby**
 - Check the box to enable **Cache Mirroring** (or uncheck to disable)
5. Click the **Save** button.

LOCKING OR UNLOCKING THE SUBSYSTEM

The lock prevents other sessions (including sessions with the same user) from making a configuration change to the controller until the lock expires or a forced unlock is done. When the user who locked the controller logs out, the lock is automatically released.

SETTING THE LOCK

To set the lock:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. In the Lock Time field, type a lock time in minutes.

1440 minutes = 24 hours

5. Click the **Lock** button.

RESETTING THE LOCK

To reset the lock with a new time:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. In the Lock Time field, type a new lock time in minutes.

1440 minutes = 24 hours

5. Click the Lock button.

RELEASING THE LOCK

To release a lock that you set:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. Click the **Unlock** button.

RELEASING A LOCK SET BY ANOTHER USER

To release somebody else’s lock:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. Check the **Force Unlock** box.
5. Click the **Unlock** button.

RESTORING FACTORY DEFAULT SETTINGS

This feature restores settings to their default values.



Caution

Use this feature only when required and only on the settings that you must reset to default in order to set them correctly.

To restore all settings to their default values:

1. Click the **Administration** tab.
2. Click the **Restore Factory Default** icon.
3. In the Restore factory default settings screen, check the boxes beside the settings you want to reset to default value:

Firmware Factory Default Settings

- Background activity settings
- Controller settings
- Enclosure settings
- FC port settings
- iSCSI port settings
- Management network settings
- Physical drive settings
- Subsystem settings

Software Factory Default Settings

- BGA scheduler settings
- Service settings
- Webserver settings
- SNMP settings
- Telnet settings
- SSH settings
- Email settings
- Netsend settings
- NTP settings
- User settings
- UPS settings

4. Click the Submit button.
5. In the Confirmation box, type the word “confirm” in the field provided and click the Confirm button.

CLEARING STATISTICS

This function clears statistical data on the RAID controllers, Fibre Channel ports, physical drives, and logical drives.

To clear subsystem statistics:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Clear Statistics** button.
4. Type the word “**confirm**” in the field provided.
5. Click the **Confirm** button.

SAVING A SERVICE REPORT

A Service Report is a detailed report covering the configuration and status of all components in your RAID system. A support technician or field engineer might request a service report for the purpose of diagnosis and troubleshooting.

To save a system configuration file:

1. Click **Save Service Report** in the Header (very top of the web interface, next to the **Help** link).

Information for the report is gathered and compiled. This action takes up to a few minutes, depending on the size of your RAID system

2. Click the **Save File** option, then click the Save button.

The report saves to your Host PC as a compressed HTML file.

3. Double-click the downloaded file to decompress it.
4. Double-click the report to open it in your default browser.

The Service Report includes the following topics:

- About – Report utility
- Battery Info – Cache backup batteries
- BBM Info – Bad Block Manager
- BGA Summary – Status and settings

The Service Report includes the following topics, continued:

- | | |
|--|---|
| • BGA Schedules – Scheduled activities | • Error Table Info – Read check, write check, and inconsistent blocks |
| • Buzzer Info | • Event Info – NVRAM – List of NVRAM events |
| • Controller Info | • Event Info – Runtime – List of Runtime events |
| • Debug Syslog – Diagnostic information | • FC Node Info |
| • Disk Array Info – ID, alias, and capacities only | • FC Device Info |
| • Disk Array Dump Info – Diagnostic information | • FC Initiator Info |
| • Disk Array Verbose Info – All disk array information | • FC Port Info |
| • Enclosure Info | • FC SFP Info |

The Service Report includes the following topics, continued:

- FC Stats Info
- Flash Image Version Info
- iSCSI Info
- LDAP Info
- LogDrive Info – Basic logical drive information
- LogDrive Dump Info – Diagnostic information
- Logical Drive Verbose Info – Full logical drive information
- Lunmap Info – LUN map type, LUN masking status, and LUN entries
- Network Info – Virtual port
- Network Maintenance Info – Maintenance mode ports
- Phydriv Info – Basic physical drive information
- Phydriv Verbose Info – Full physical drive information
- PD SMART Info – Physical drive ID, model, type, and SMART status
- PSU Wattage Info – Enclosure power consumption, power supply input and output, and power on time
- SWMGT Info – Software management
- Service Setting – Email
- Service Setting – Netsend
- Service Setting – NTP
- Service Setting – SLP
- Service Setting – SNMP
- Service Setting – SSH
- Service Setting – Telnet
- Service Setting – Webserver
- Sessions Info
- Spare Info – Basic spare drive information
- Spare Dump Info – Diagnostic information
- Spare Verbose Info – Full spare drive information
- Statistic Info
- Subsystem info
- UPS Info
- User Info

IMPORTING A CONFIGURATION SCRIPT

You can write a CLI configuration script to automatically configure your Vess R2600 subsystem. The script must be a plain, non-encrypted text file. From there, you can import the script from the Host PC and perform the configuration automatically.



Cautions

Do NOT attempt to write or modify a configuration script until you receive guidance from Technical Support. See "Frequently Asked Questions" on page 225.

Importing a configuration script overwrites the current settings on your Vess R2600 subsystem.

Or you can save the configuration from one Vess R2600 RAID subsystem, export it, and then import it to automatically configure your other Vess R2600 RAID subsystems. To import a configuration script:

1. Click the **Administration** tab.
2. Click the **Import/Export** icon.
3. Click the **Import** option.
4. Choose **Configuration Script** from the **Type** drop-down menu.
5. Click the **Browse** button and navigate to the configuration script and click the **OK** button.
6. Click the **Next** button.

- The system verifies that the file is a valid configuration script and displays any errors or warnings.
7. Click the **Submit** button to continue.
 8. In the **Confirmation** box, type the word "**confirm**" in the field provided and click the **Confirm** button.

The configuration script is imported and applied automatically.

EXPORTING A CONFIGURATION SCRIPT

You can save the configuration from one Vess R2600 RAID subsystem, export it, and then import it to automatically configure your other Vess R2600 RAID subsystems.

To export a configuration script:

1. Click the **Administration** tab.
2. Click the **Import/Export** icon.
3. Click the **Export** option.

4. Choose **Configuration Script** from the **Type** drop-down menu.
5. Click the **Submit** button.
6. In the Open dialog box, click the **Save File** option, then click the **OK** button.

The file is saved to your PC as "Configscript.txt".



Cautions

Do NOT attempt to write or modify a configuration script until you receive guidance from Technical Support. See "Frequently Asked Questions" on page 225.

RESTARTING THE SUBSYSTEM

This function shuts down the subsystem and then restarts it.

To restart the subsystem:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Choose the option to apply the restart to the **Subsystem**, **Controller 1 only** or **Controller 2 only**.
4. Click the **Shutdown/Restart** button.
5. Click the **Restart** button.
6. Type the word "confirm" in the field provided.
7. Click the **Confirm** button.

When the controller shuts down, your WebPAM PROe connection is lost.

8. Wait at least two minutes.
9. In your browser, click **Logout** in the WebPAM PROe Header, then log in again.

If you cannot log in immediately, wait 30 seconds and try again.

SHUTTING DOWN THE SUBSYSTEM

This function shuts down the RAID subsystem without restarting it.

To shutdown the subsystem:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Choose the option to apply the shutdown to the **Subsystem**, **Controller 1 only** or **Controller 2 only**
4. Click the **Shutdown/Restart** button.
5. Click the **Shutdown** button.
6. Type the word “**confirm**” in the field provided.
7. Click the **Confirm** button.

When the controller shuts down, your WebPAM PROe connection is lost.

8. Wait at least two minutes.



Important

If your RAID subsystem manages JBOD expansion units, you must follow the proper startup procedure.

RESTARTING THE SUBSYSTEM AFTER A SHUTDOWN

To start the RAID subsystem:

1. Press the Power button on the front left side of the device being restarted.
2. Wait at least two minutes.
3. Open your browser and log into WebPAM PROe.

If you cannot log in immediately, wait 30 seconds and try again.

MANAGING RAID CONTROLLERS

RAID controller management includes:

- "Viewing Controller Information" on page 58
- "Making Controller Settings" on page 59
- "Locating a Controller" on page 60
- "Viewing the Flash Image Information" on page 60
- "Viewing Battery Information" on page 62
- "Reconditioning a Battery" on page 62
- "Silencing the Buzzer" on page 63

VIEWING CONTROLLER INFORMATION

To view controller information:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **View** button.

Controller information includes:

- | | |
|----------------------------|----------------------------------|
| • Controller ID | • Alias – If assigned * |
| • Readiness Status | • Operational Status |
| • Power On Time | • SCSI Protocol Supported |
| • LUN Mapping method | • Part Number |
| • Serial Number | • WWN – World Wide Name |
| • Hardware Revision | • Dirty Cache Usage – Percentage |
| • Cache Usage – Percentage | • Firmware Version |
| • Boot Loader Version | • Software Version |
| • Firmware Build Date | |
| • Software Build Date | |

4. Click the **Advanced Information** tab.

Advanced controller information includes:

- Slot 1 Memory Type
- Slot 2 Memory Type
- LUN Affinity *
- Controller Role
- Flash Size
- NVRAM Size
- Coercion *
- SMART *
- Write Back Cache Flush Interval *
- Adaptive Writeback Cache *
- Forced Read Ahead (cache) *
- Power Saving Standby Time *
- Cache Line Size
- Backup Flash Status
- Slot 1 Memory Size
- Slot 2 Memory Size
- ALUA *
- Flash Type
- NVRAM Type
- Preferred Cache Line Size
- Coercion Method *
- SMART Polling Interval *
- Enclosure Polling Interval *
- Host Cache Flushing *
- Power Saving Idle Time *
- Power Saving Stopped Time *
- Advanced Battery Flash Backup Enabled *
- Backup Flash Size

Items with an asterisk (*) are adjustable under Controller Settings.

MAKING CONTROLLER SETTINGS

In a dual-controller RAID subsystem, settings made to one controller are applied to both controllers.

To make controller settings:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **Settings** button.

4. Make settings changes as required:

- Enter, change or delete the alias in the **Alias** field.
- **LUN Affinity** – Choose an enable/disable option from the drop-down menu.
RAID controllers must be set to Active-Active. See "Making Subsystem Settings" on page 54 and "LUN Affinity" on page 194.
- **ALUA** – Choose an enable/disable option from the drop-down menu.
RAID controllers must be set to Active-Active. See Making Subsystem Settings and “ALUA” on page 251.
- **SMART Log** – Check the box to enable or uncheck to disable.
- **SMART Polling Interval** – **Enter** a value into the field, 1 to 1440 minutes
- **HDD Power Saving** – Choose time periods from the drop-down menus.
After an HDD has been idle for the set period of time:
 Power Saving Idle Time – Parks the read/write heads.
 Power Saving Standby Time – Lowers disk rotation speed.
 Power Saving Stopped Time – Spins down the disk (stops rotation).
- **Coercion** – Check the box to enable or uncheck to disable.
- **Coercion Method** – Choose a method from the drop-down menu:
 GBTruncate
 10GBTruncate
 GrpRounding
 TableRounding
- **Write Back Cache Flush Interval** – **Enter** a value into the field, 1 to 12 seconds.
- **Enclosure Polling Interval** – 15 to 255 seconds.
- **Adaptive Writeback Cache** – Check the box to enable or uncheck to disable. See "Adaptive Writeback Cache" on page 195.
- **Host Cache Flushing** – Check the box to enable or uncheck to disable. See "Host Cache Flushing" on page 195.
- **Forced Read Ahead (cache)** – Check the box to enable or uncheck to disable. See "Forced Read-Ahead Cache" on page 195.
- **Advanced Battery Flash Backup** - Check the box to enable or uncheck to disable. See "Advanced Battery Flash Backup" on page 196.

5. Click the **Save** button.



Notes

Power Management must be enabled on the disk array for the HDD Power Saving settings to be effective. See "Making Disk Array Settings" on page 93

Power Management functions are limited to the features your HDDs actually support.

VIEWING CONTROLLER STATISTICS

To view controller statistics:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **View** button.
4. Click the **Statistics** tab.

Controller statistics include:

- Data Transferred
- Read Data Transferred
- Errors
- Read Errors
- I/O Requests
- Read IO Requests
- Statistics Start Time
- Write Data Transferred
- Non-Read/Write Errors
- Write Errors
- Non-Read/Write Requests
- Write I/O Requests
- Statistics Collection Time



Note

To clear controller statistics, see "Clearing Statistics" on page 55.

LOCATING A CONTROLLER

This feature causes the controller LEDs to blink for one minute to assist you in locating the controller on a RAID subsystem or JBOD expansion unit.

To locate a controller:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **Locate** button.

The controller status LEDs blink for one minute.

VIEWING THE FLASH IMAGE INFORMATION

To view the flash image information for the RAID subsystem enclosure:

1. Click the **Administration** tab.
2. Click the **Image Version** icon.
3. Click the **Enclosure** you want to see and click the triangular button.

RAID subsystems have the following components in their flash image:

- Kernel
- Firmware
- Software
- Ramdisk
- SEP Firmware
- OEM Customization
- BIOS
- 6G Expander
- System Libraries
- Applications
- Mount Scripts
- PLX EEPROM Image

- **Running** – The version that is currently running on the subsystem or expansion unit.
- **Flashed** – This version was updated but does not run until the subsystem restarts.

See Updating Firmware on a RAID Subsystem.

JBOD expansion units have only one component in their flash image, SEP firmware. It only appears as running.

UPDATING FIRMWARE ON A RAID SUBSYSTEM

Use this function to flash (update) the firmware on the Vess R2600.

Download the latest firmware image file from PROMISE support:

<http://www.promise.com/support/> and save it to your Host PC or TFTP server.



Important

Verify that no background activities are running on the RAID subsystem.

To update the firmware on the RAID subsystem and JBOD expansion units:

1. Click the **Administration** tab.
2. Click the Firmware **Update** icon.
3. Click the **Controller Firmware Update** tab.

The Controller Firmware Update screen appears showing the current Image Version Number and Build Date.

4. Choose a download option:
 - **Local File through HTTP** – Click the **Browse** button, locate the firmware image file, click the file to choose it, then click the **Open** button.
 - **TFTP Server** – **Enter** the TFTP Server host name or IP address, port number and file name.

5. Optional. Check the Non-disruptive Image Update (NDIU) box.

NDIU updates the RAID controllers and I/O modules one at a time, enabling I/O operations continue during the firmware update. Updates with this option take a longer period of time to complete. Only Vess R2600 models support this feature.

6. Click the **Next** button.

The next screen shows the Flash Image (firmware image file) Version Number and Build Date.

7. Click the **Submit** button.

The progress of the update displays.



Warning

Do NOT power off the RAID subsystem during the update!

Do NOT move to any other screen until the firmware update operation is completed!

When the update is completed a message tells you to reboot the subsystem,

8. Click the **OK** button.
 - If you chose the Disruptive Flash Method, the RAID subsystem and JBOD expansion units automatically restart.
 - If you chose the Non-Disruptive Flash Method, the system automatically flashes and restarts the RAID controllers one at a time.

AUTOMATIC RESTART

If you did NOT check the NDIU box, the RAID subsystem and JBOD expansion units automatically restart. That action temporarily disrupts I/O operations and drops your WebPAM PROe connection.

To reestablish your WebPAM PROe connection:

1. Wait no less than two minutes.
2. Click **Logout** in the WebPAM PROe Header, then log in again.

If you cannot log in, wait 30 seconds and try again.
3. In your browser, click Logout in the WebPAM PROe Header, then log in again.

If you cannot log in immediately, wait 30 seconds and try again.

VIEWING BATTERY INFORMATION

Batteries maintain power to the controller cache in the event of a power failure, thus protecting any data that has not been written to a physical drive.

To view battery information:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the battery you want, then click the **View** button.

Battery information includes:

- Battery ID
- Operational status – Fully charged, recondition means a reconditioning is in process
- Battery chemistry – LiON, etc.
- Remaining capacity – Battery capacity as a percentage
- Battery cell type – Number of cells
- Estimated hold time – Time in hours that the battery can power the cache
- Temperature threshold discharge – Maximum temperature allowed when the battery is discharging
- Temperature threshold charge – Maximum temperature allowed when the battery is charging
- Battery temperature – Actual battery temperature
- Cycle count – Number of times the battery was reconditioned
- Voltage in millivolts
- Current in milliamps

RECONDITIONING A BATTERY

Batteries maintain power to the controller cache in the event of a power failure, thus protecting any data that has not been written to a physical drive. Reconditioning is the action of discharging and recharging a battery to preserve its capacity and performance.

Reconditioning is a background activity, it might affect I/O performance. When the recondition is completed, the battery’s cycle count increments by one.

Battery reconditioning is disabled by default. You can change the reconditioning status and schedule.

To recondition a battery immediately:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the battery you want, then click the **Recondition** button.

Battery operations status changes to “Recondition” and the battery’s remaining capacity and estimated hold time fall and rise reflecting the discharge and recharge cycles of the reconditioning. That behavior is normal.

MAKING SCHEDULE CHANGES

To make changes the scheduled battery reconditioning:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of Background Activities displays.
3. Click the **Scheduler** button.
4. Mouse-over **Battery Reconditioning** and click the Settings button.
5. Make settings changes as required:
 - Start Time
 - Uncheck the Enable This Schedule box to disable this activity.
 - Recurrence Pattern
 - Start From
 - End On
6. Click the **Save** button to apply the new settings.

Buzzer Settings

To make buzzer settings:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Buzzer and click the **Settings** button.
4. Check the **Enable Buzzer** box to enable the buzzer, or uncheck the box to disable.
5. Click the **Save** button.

Silencing the Buzzer



Caution

This action disables the buzzer for all events.

To silence the buzzer:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Buzzer and click the **Settings** button.
4. Uncheck the **Enable Buzzer** box.
5. Click the **Save** button.



Note

The alarm can also be muted by pressing the Mute Alarm button on the front left side of the Vess R2600 unit. See "Figure 2 - 23: Left side of the front of the Vess R2600" on page 31.

Managing Enclosures

Enclosure management includes the following functions:

- "Viewing Enclosure Topology" on page 64
- "Viewing the Enclosures Summary" on page 65
- "Making Enclosure Settings" on page 65
- "Locating an Enclosure" on page 65
- "Viewing FRU VPD Information" on page 66
- "Viewing Power Supply Status" on page 66
- "Viewing Cooling Unit Status" on page 66
- "Viewing Temperature Sensor Status" on page 66
- "Viewing Voltage Sensor Status" on page 66

VIEWING ENCLOSURE TOPOLOGY

This feature displays the connection topology of the Vess R2600 subsystem. Topology refers to the manner in which the data paths among the enclosures are connected. There are three methods:

- **Individual Subsystem** – A single subsystem
- **JBOD Expansion** – Managed through one subsystem or head unit
- **RAID Subsystem Cascading** – Managed through one subsystem or head unit


For more information about connections, see "Making Management and Data Connections" on page 18.

To view enclosure topology:

1. Click the **Device** tab.
2. Click the **Topology** icon.

The topology or data connections of your system displays.

Figure 4 - 5: Enclosure Topology menu



VIEWING THE ENCLOSURES SUMMARY

Enclosure Management includes information, status, settings and location. To access Enclosure Management:

1. Click the **Device** tab.
2. Click the **Component List** icon.

The following information is shown:

- Enclosure ID number
- Status
- Enclosure Type
- Status Description (specific components in need of attention, if any)

LOCATING AN ENCLOSURE

To locate an enclosure:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure you want, then click the **Locate** button.

The enclosure LEDs blink for one minute.

VIEWING ENCLOSURE INFORMATION

To view enclosure information:

1. Click the **Device** tab.
2. Click the **Component List** icon.

3. Click the Enclosure and click the **View** button.

Enclosure information includes:

- Enclosure ID
- Enclosure Type
- Enclosure Warning Temperature Threshold
- Controller Warning Temperature Threshold
- Max Number of Controllers
- Max Number of Fans
- Max Number of Temperature Sensors
- Max Number of Batteries
- Enclosure Critical Temperature Threshold
- Controller Critical Temperature Threshold
- Max Number of Physical Drive Slots
- Max Number of Blowers
- Max Number of Power Supply Units
- Max Number of Voltage Sensors

For information on Enclosure problems, see "Diagnosing an Enclosure Problem" on page 209.

MAKING ENCLOSURE SETTINGS

To make Enclosure settings:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **Settings** button.

Enclosure settings include:

- Enclosure Warning Temperature Threshold [51-55]°C
- Enclosure Critical Temperature Threshold [56-68]°C
- Controller Warning Temperature Threshold [71-75]°C
- Controller Critical Temperature Threshold [76-80]°C

4. In the field provided, type the temperature in degrees C for each threshold value.
5. Click the **Save** button.

VIEWING FRU VPD INFORMATION

FRU VPD refers to Vital Product Data (VPD) information about Field Replaceable Units (FRU) in the enclosure. The number and type of FRU depends on the subsystem model.

To view FRU VPD information:

To make Enclosure settings:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **FRU VPD** button.

Use this information when communicating with Technical Support and when ordering replacement units.

For contact information, see "Contacting Technical Support" on page 228.

VIEWING POWER SUPPLY STATUS

To view the status of the power supplies and the fans that cool those power supplies:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.
4. Scroll down to view the power supplies.

The screen displays the operational and fan status of the power supplies. If any status differs from normal or the fan speed is below the Healthy Threshold value, a malfunction is indicated in the Status column.

See "Replacing a Power Supply" on page 177.

VIEWING COOLING UNIT STATUS

To view the status of the cooling units:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.

4. Scroll down to view the Blowers.

The screen displays the status and speed of the cooling units. If blower speed is below the Healthy Threshold, a malfunction is indicated in the Status column. See "Diagnosing an Enclosure Problem" on page 209.

VIEWING TEMPERATURE SENSOR STATUS

To view the status of the temperature sensors:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.
4. Scroll down to view the Temperature Sensors.

If any temperature exceeds the Healthy Threshold value, an overheat condition exists in the enclosure.

See "Making Enclosure Settings" and "Diagnosing an Enclosure Problem" on page 209.

VIEWING VOLTAGE SENSOR STATUS

To view the status of the voltage sensors:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.
4. Scroll down to view the Voltage Sensors.

If any voltage is outside the Healthy Threshold values, a voltage malfunction in the enclosure is indicated in the Status column. See "Diagnosing an Enclosure Problem" on page 209.

MANAGING UPS UNITS

Uninterruptible Power Supply (UPS) Management includes the following functions:

- "Viewing a List of UPS Units" on page 67
- "Viewing UPS Information" on page 68

VIEWING A LIST OF UPS UNITS

To view a list of UPS units supporting the Vess R2600:

1. Click the Device tab.
2. Click the UPS icon.

Information in the UPS List includes:

- **ID** – The ID number of the UPS
- **Status** – OK means Normal.
On AC means the UPS is connected to a viable external AC power source.
On Battery means the external AC power source is offline and the UPS is running on battery power.
- **Model** – Model name of the UPS
- **Battery Capacity** – Backup capacity expressed as a percentage.
- **Loading Ratio** – Actual output of UPS as a percentage of the rated output. See the Note below.
- **Remaining Minutes** – Number of minutes the UPS is expected to power your system in the event of a power failure.



Note

The maximum recommended Loading Ratio varies among models of UPS units. The general range is 60% to 80%. If the reported Loading Ratio exceeds the recommended value for your UPS unit:

Have fewer subsystems or peripherals connected to this UPS unit.

Add more UPS units, or use a higher-capacity UPS unit, to protect your RAID systems.

MAKING UPS SETTINGS

These settings control how the Vess R2600 subsystem detects the UPS unit and responds to data reported by the UPS unit.

To make UPS settings:

1. Click the **Device** tab.
2. Click the **UPS** icon.
3. Click the **UPS Settings** button.
4. Perform the following actions as required:
 - Verify the Current UPS Communication method. See Note 1:
SNMP – Network connection.
USB
Unknown – No connection.
 - Choose a Detection Setting from the drop-down menu:
Automatic – Default. If a UPS is detected when the subsystem boots, the settings changes to Enable.
Enable – Monitors UPS. Settings changes, reports warnings, and logs events.
Disable – Does not monitor UPS.
 - Type values into the Threshold fields. See Note 2:
Running Time Remaining Threshold – Actual time below this value resets adaptive writeback cache to writethrough.
Warning Temperature Threshold – Actual temperature above this value triggers a warning and logs an event.
Loading Ratio Threshold – Actual loading ratio (percentage) above this threshold triggers a warning and logs an event. See Note 3.
 - For UPS units with network cards, type the IP addresses or DNS names in fields UPS 1 and UPS 2. See Note 4.
5. Press **Submit** to save your settings.

Note 1: Vess R2600 supports multiple UPS units using network or USB connections, but not a combination of both methods.

Note 2: Detection Setting must be set to Auto. If a UPS is detected, the settings changes to Enable.

Note 3: The maximum recommended Loading Ratio varies among models of UPS units. The general range is 60% to 80%.

Note 4: To specify UPS units by DNS names, ask your IT administrator to add the DNS names to the DNS server, before you make UPS settings.

VIEWING UPS INFORMATION

To view information about a specific UPS unit:

1. Click the **Device** tab.
2. Click the **UPS** icon.
3. Mouse-over UPS and click the **View** button.

UPS information includes:

- **UPS ID**
- **Model Name**
- **Serial Number**
- **Firmware Version**
- **Manufacture Date**
- **Voltage Rating** – Output voltage of the UPS.
- **Battery Capacity** – Backup capacity expressed as a percentage.
- **Remaining Backup Time** – Number of minutes the UPS is expected to power your system in the event of a power failure.
- **Loading Ratio** – Actual output of UPS as a percentage of the rated output. See the Note below.
- **Temperature** – Reported temperature of the UPS unit.



Note

The maximum recommended Loading Ratio varies among models of UPS units. The general range is 60% to 80%. If the reported Loading Ratio exceeds the recommended value for your UPS unit:
Have fewer subsystems or peripherals connected to this UPS unit.
Add more UPS units, or use a higher-capacity UPS unit, to protect your RAID systems.

MANAGING NETWORK CONNECTIONS

Network Connections Management includes the following functions:

- "Making Virtual Management Port Settings" on page 68
- "Making Maintenance Mode Settings" on page 69

MAKING VIRTUAL MANAGEMENT PORT SETTINGS

The Vess R2600 subsystem has a virtual management port, enabling you to log into a Vess R2600 with dual controllers using one IP address.

Before you change settings, please see "About IP Addresses" on page 36.

You initially made these settings during subsystem setup. You can change them later as required.



Caution

Changing virtual management port settings can interrupt your WebPAM PROe connection and require you to log in again.

To make virtual management port settings:

1. Click the **Administration** tab.
2. Click the **Network Management** icon.
3. Click the **Virtual Management Port** tab.
4. Click the protocol family whose settings you want to change and click the **Configuration** button.
5. Make the following settings are needed:
 - Check the **Enable** box to enable this protocol family.
 - Check the **Enable DHCP** box to enable a DHCP server to make your network settings. DHCP is currently supported in IPv4 only.
 - For manual network settings, type the RAID subsystem's IP address, subnet mask, gateway IP address, and DNS server IP address into the fields provided.
6. Click the **Submit** button.

MAKING MAINTENANCE MODE SETTINGS

Each controller has its own IP addresses for access when the controller goes into maintenance mode. For more information, see "Maintenance Mode" on page 211.

Before you change settings, please see “About IP Addresses” on page 36.

To make maintenance mode settings:

1. Click the **Administration** tab.
2. Click the **Network Management** icon.
3. Click the **Maintenance Mode** tab.
4. Click the controller and protocol family whose settings you want to change and click the **Configuration** button.
5. Make the following settings are needed:
 - Check the **Enable** box to enable this protocol family.
 - Check the **Enable DHCP** box to enable a DHCP server to make your network settings. DHCP is currently supported in IPv4 only.
 - For manual network settings, type the **IP address, subnet mask, gateway IP address**, and **DNS server IP address** into the fields provided.
6. Click the **Sub mit** button.

MANAGING USERS

User management includes:

- "Viewing User Information" on page 69
- "Changing User Passwords" on page 70
- "Importing a User Database" on page 71
- "Exporting a User Database" on page 72

The **Administrator** or a **Super User** can perform these tasks.

VIEWING USER INFORMATION

To view user information:

1. Click the **Administration** tab.
2. Click the **User Management** icon.

The list of users displays. User information includes:

- User name
- Status
- Privilege level
- Display name
- Email address

CREATING A USER

This action requires **Administrator** or **Super User** privileges.

To create a user:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. Click the **Add User** button.
4. In the **Add User** dialog box, enter the information in the fields provided:
 - Name – This is the user’s login name
 - Display Name
 - Password
 - Retype Password
 - User Email – Required for event notification
5. Choose a privilege level from the drop-down menu. See the table below for a description of the privilege types.
6. (Optional) Uncheck the Enable box to disable this User account.
7. Click the **Save** button. The user is added to the list.



Important

For this user to receive event notification, Click the new user and click the **Subscription** button.

User Privileges	
Level	Meaning
View	Allows the user to see all status and settings but not to make any changes
Maintenance	Allows the user to perform maintenance tasks including Rebuilding, PDM, Media Patrol, and Redundancy Check
Power	Allows the user to create (but not delete) disk arrays and logical drives, change RAID levels, change stripe size; change settings of components such as disk arrays, logical drives, physical drives, and the controller
Super	Allows the user full access to all functions including create and delete users and changing the settings of other users, and delete disk arrays and logical drives. The default “administrator” account is a Super User

MAKING USER SETTINGS

This action requires **Administrator** or a **Super User** privileges.

To make user settings:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click **Settings**.
4. Make settings changes as required:
 - For the **Enable** box, check to enable this user account, uncheck to disable this user account
 - In the User Settings dialog box, enter a new **Display Name** or **User Email** address
 - Choose a new **Privilege** level from the drop-down menu. See the table on the next page.
5. Click the **Save** button.

CHANGING USER PASSWORDS

This action requires **Administrator** or **Super User** privileges.

To change a user’s password:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click **Change Password**.
4. In the Change Password dialog box, enter the information in the fields provided:
 - New Password
 - Retype Password
5. Click the **Save** button.

DELETING A USER

This action requires **Administrator** or **Super User** privileges

.



Note

You cannot delete the Administrator.

To delete a user:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click the **Delete** button.
4. In the **Confirmation** box, type the word “**confirm**” in the field provided and click the **Confirm** button.

SETTING USER EVENT SUBSCRIPTIONS

By default, all users have event notification:

- Enabled
- Set to the Major (severity) level for all events

Subscribing users receive notification of events at the chosen severity level and all higher levels.



Note

Each user must have a valid Email address to receive events. See “Making User Settings” below.

Changing a user subscription requires **Administrator** or **Super User** privileges.

To set a user event subscription:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click the **Subscription** button.
4. Make settings changes as required:
 - For the **Enable Event Notification** box, check to enable for this user, uncheck to disable.
 - Click to change the priority options for each category of event.
5. Click the **Save** button.

IMPORTING A USER DATABASE

You can save the user information and settings from one Vess R2600 RAID subsystem, export it, and then import it to automatically configure your other Vess R2600 RAID subsystems.



Caution

Importing a user database overwrites the current users and user settings on your Vess R2600 subsystem.

To import a user database:

1. Click the **Administration** tab.
2. Click the **Import/Export** icon.
3. Click the **Import** option.
4. Choose **User Database** from the **Type** drop-down menu.
5. Click the **Browse** button and navigate to the user database file and click the **OK** button.
6. Click the **Next** button.

The system verifies that the file is a valid user database and displays any errors or warnings.

7. Click the **Submit** button to continue.
8. In the **Confirmation** box, type the word “**confirm**” in the field provided and click the Confirm button.

The user database is imported and applied automatically.

EXPORTING A USER DATABASE

You can save the user information and settings from one Vess R2600 RAID subsystem, export it, and then import it to automatically configure your other Vess R2600 RAID subsystems.

To export a user database:

1. Click the **Administration** tab.
2. Click the **Import/Export** icon.
3. Click the **Export** option.
4. Choose **User Database** from the **Type** drop-down menu.
5. Click the **Submit** button.
6. In the **Open** dialog box, click the **Save File** option, then click the **OK** button.

The file is saved to your PC as “User.dat”.



Note

The user database file is not designed to be opened or edited in the field..

MANAGING BACKGROUND ACTIVITIES

Background activity management includes:

- "Viewing Current Background Activities" on page 73
- "Viewing Scheduled Background Activities" on page 73
- "Adding a Scheduled Background Activity" on page 73
- "Changing a Background Activity Schedule" on page 74
- "Enabling or Disabling a Scheduled Background Activity" on page 74
- "Deleting a Scheduled Background Activity" on page 74
- "Media Patrol" on page 75
- "Redundancy Check" on page 75
- "Initialization" on page 75
- "Rebuild" on page 76
- "Migration" on page 76
- "PDM" on page 76
- "Transition" on page 77
- "Synchronization" on page 77
- "Battery Reconditioning" on page 77
- "Running Spare Check" on page 103

Background activities perform a variety of preventive and remedial functions on your physical drives, disk arrays, logical drives, and other components.

You can run a background activity immediately or schedule it to run at a later time. Scheduling options are described below.

Setting options for each activity are listed after the scheduling options. These settings determine how the background activity affects I/O performance.

VIEWING CURRENT BACKGROUND ACTIVITIES

To view a list of current background activities:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

Currently running activities show a progress bar.

VIEWING SCHEDULED BACKGROUND ACTIVITIES

To view a list of scheduled background activities:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Click the **Scheduler** button.

The list of currently scheduled background activities appears.

ADDING A SCHEDULED BACKGROUND ACTIVITY

To add a new scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Click the **Scheduler** button.

The list of currently scheduled background activities appears.

4. Click the **Add Schedule** button.
5. Check the **Enable Media Patrol** box to enable, uncheck to disable.

This settings enables or disables Media Patrol for your entire RAID system.

6. Click the **Confirm** button.
7. Choose the option for the activity you want:
 - Media Patrol
 - Redundancy Check
 - Spare Check
 - Battery Recondition

8. Choose a **Start Time** from the drop-down menus.

The menus have a 24-hour clock.

9. Choose a **Recurrence Pattern** option, daily, weekly, or monthly.
 - For the Daily option, enter an interval in the Every field.
 - For the Weekly option, enter an interval in the Every field and choose one or more days of the week.
 - For the Monthly option, choose, Day of the Month option then choose a number from the drop-down menu. The day of the week option then choose the day of the month from the drop-down menus.

10. Choose a **Start From** date from the drop-down menus.

11. Choose an **End On** option,
 - No end date or perpetual.
 - End after a specific number of activity actions.
 - Until date from the drop-down menus.

12. For **Redundancy Check**, choose,
 - **Auto Fix** option – Attempts to repair the problem when it finds an error. Check to enable
 - **Pause on Error** option – The process stops when it finds a non-repairable error. Check to enable
 - **Select LD** – Check the boxes for the logical drives to run Redundancy Check. Check at least one logical drive

13. Click the **Save** button.

CHANGING A BACKGROUND ACTIVITY SCHEDULE

To change an existing scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Click the **Scheduler** button.

The list of currently scheduled background activities appears.

4. Click the background activity and click the **Settings** button.
5. Make settings changes as required:
 - Choose a **Start Time** from the drop-down menus.
The menus have a 24-hour clock.
 - Choose a **Recurrence Pattern** option, daily, weekly, or monthly.
For the Daily option, enter an interval in the Every field.
For the Weekly option, enter an interval in the Every field and choose one or more days of the week.
For the Monthly option, choose the Day of the Month option or the day of the week option, and choose the day from the drop-down menu.
 - Choose a **Start From** date from the drop-down menus.
 - Choose an **End On** option,
No end date or perpetual.
End after a specific number of activity actions.
Until date from the drop-down menus.
 - For **Redundancy Check**, choose,
Auto Fix option – Attempts to repair the problem when it finds an error. Check to enable
Pause on Error option – The process stops when it finds a non-repairable error. Check to enable
Select LD – Check the boxes for the logical drives to run Redundancy Check. Check at least one logical drive
6. Click the **Save** button.

ENABLING OR DISABLING A SCHEDULED BACKGROUND ACTIVITY

Background activity schedules are enabled by default when you create the schedule. If you want to stop a background activity now but plan to use it again in the future, disable the scheduled activity rather than deleting it.

To enable or disable change an existing scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Click the **Scheduler** button.

The list of currently scheduled background activities appears.

4. Click the background activity and click the **Settings** button.
5. Uncheck the **Enable This Schedule** box to disable this schedule.

Check the box to enable this schedule.

6. Click the **Save** button.

DELETING A SCHEDULED BACKGROUND ACTIVITY

To change an existing scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Click the **Scheduler** button.

The list of currently scheduled background activities appears.

4. Click the background activity and click the **Delete** button.
5. In the confirmation box, click the confirm button.

MEDIA PATROL

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. Media Patrol checks are enabled by default on all disk arrays and spare drives. Media Patrol is concerned with the media itself, not the data recorded on the media. If Media Patrol encounters a critical error, it triggers PDM if PDM is enabled on the disk array.

See "Making Disk Array Settings" on page 93, and "Media Patrol" on page 75.

MAKING MEDIA PATROL SETTINGS

To make Media Patrol settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Click the **Settings** button.
4. Check the **Enable Media Patrol** box to enable, uncheck to disable.

This settings enables or disables **Media Patrol** for your entire RAID system.

5. Click the **Confirm** button.

You can also enable or disable **Media Patrol** on individual disk arrays.

REDUNDANCY CHECK

Redundancy Check is a routine maintenance procedure for fault-tolerant disk arrays (those with redundancy) that ensures all the data matches exactly. Redundancy Check can also correct inconsistencies.

See "Redundancy Check on a Logical Drive" on page 99

MAKING REDUNDANCY CHECK SETTINGS

To make Redundancy Check settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Click the **Settings** button.
4. Click the **Redundancy Check Rate** drop-down menu and choose a rate:
 - **Low** – Fewer system resources to Redundancy Check, more to data read/write operations.
 - **Medium** – Balances system resources between Redundancy Check and data read/write operations.
 - **High** – More system resources to Redundancy Check, fewer to data read/write operations.
5. Click the **Confirm** button.

INITIALIZATION

Technically speaking, **Initialization** is a foreground activity, as you cannot access a logical drive while it is initiating.

Initialization is normally done to logical drives after they are created from a disk array. Initialization sets all data bits in the logical drive to zero. The action is useful because there may be residual data on the logical drives left behind from earlier configurations. For this reason, Initialization is recommended whenever you create a logical drive.

See "Initializing a Logical Drive" on page 99.

MAKING INITIALIZATION SETTINGS

To make initialization settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

- Click the **Settings** button.
- Click the Logical Drive Initialization Rate drop-down menu and choose a rate:
 - Low** – Fewer system resources to Initialization, more to data read/write operations.
 - Medium** – Balances system resources between Initialization and data read/write operations.
 - High** – More system resources to Initialization, fewer to data read/write operations.
- Click the **Confirm** button.

REBUILD

When you rebuild a disk array, you are actually rebuilding the data on one physical drive.

- When a physical drive in a disk array fails and a spare drive of adequate capacity is available, the disk array begins to rebuild automatically using the spare drive.
- If there is no spare drive of adequate capacity, but the Auto Rebuild function is ENABLED, the disk array begins to rebuild automatically as soon as you remove the failed physical drive and install an unconfigured physical drive in the same slot. See "Making Rebuild Settings" on page 76.
- If there is no spare drive of adequate capacity and the Auto Rebuild function is DISABLED, you must replace the failed drive with an unconfigured physical drive, then perform a **Manual Rebuild**.

See "Rebuilding a Disk Array" on page 95 and "Spare Drives" on page 193.

Also see "Disk Array Degraded / Logical Drive Critical" on page 213 and "Disk Array Offline / Logical Drive Offline" on page 214.

MAKING REBUILD SETTINGS

- Click the **Administration** tab.
- Click the **Background Activities** icon.

The list of background activities appears.
- Click the **Settings** button.
- Click the **Rebuild Rate** drop-down menu and choose a rate:
 - Low** – Fewer system resources to the Rebuild, more to data read/write operations.
 - Medium** – Balances system resources between the Rebuild and data read/write operations.
 - High** – More system resources to the Rebuild, fewer to data read/write operations.
- Check the **Enable Auto Rebuild** box to enable Auto Rebuild (rebuilds when you swap out the failed drive with a new one).
- Click the **Confirm** button.

MIGRATION

The term “Migration” means either or both of the following:

- Change the RAID level of a logical drive.
- Expand the storage capacity of a logical drive.

See “Migrating a Logical Drive’s RAID Level” on page 169 and "RAID Level Migration" on page 189.

MAKING MIGRATION SETTINGS

To make migration settings:

- Click the **Administration** tab.
- Click the **Background Activities** icon.

The list of background activities appears.

- Click the **Settings** button.
- Click the **Migration Rate** drop-down menu and choose a rate:
 - Low** – Fewer system resources to Migration, more to data read/write operations.
 - Medium** – Balances system resources between Migration and data read/write operations.
 - High** – More system resources to Migration, fewer to data read/write operations.
- Click the **Confirm** button.

PDM

Predictive Data Migration (PDM) is the migration of data from the suspect physical drive to a spare drive, similar to rebuilding a logical drive. But unlike Rebuilding, PDM constantly monitors your physical drives and automatically copies your data to a spare drive before the physical drive fails and your logical drive goes Critical.

See "Running PDM on a Disk Array" on page 94 and "PDM" on page 184.

MAKING PDM SETTINGS

To make PDM settings:

- Click the **Administration** tab.
- Click the **Background Activities** icon.

The list of background activities appears.

- Click the **Settings** button.

4. Make the following settings are required:
- Click the **PDM Rate** drop-down menu and choose a rate:
 - Low** – Fewer system resources to PDM, more to data read/write operations.
 - Medium** – Balances system resources between PDM and data read/write operations.
 - High** – More system resources to PDM, fewer to data read/write operations.
 - Highlight the current values in the block threshold fields and input new values.
 - Reassigned block threshold range is 1 to 512 blocks.
 - Error block threshold range is 1 to 2048 blocks.
5. Click the **Confirm** button.

TRANSITION

Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare drive.

See "Running a Transition on a Spare Drive" on page 103.

MAKING TRANSITION SETTINGS

To make Transition settings:

- Click the **Administration** tab.
- Click the **Background Activities** icon.

The list of background activities appears.

- Click the **Settings** button.
- Click the **Transition Rate** drop-down menu and choose a rate:
 - Low** – Fewer system resources to Transition, more to data read/write operations.
 - Medium** – Balances system resources between Transition and data read/write operations.
 - High** – More system resources to Transition, fewer to data read/write operations.
- Click the **Confirm** button.

SYNCHRONIZATION

Synchronization is automatically applied to redundant logical drives when they are created. Synchronization recalculates the redundancy data to ensure that the working data on the physical drives is properly in sync.

Mouse-over on the logical drive, click the View button, and look under Logical Drive Information beside the line that says Synchronized. A **Yes** means the logical drive was synchronized. See "Viewing Logical Drive Information" on page 96

MAKING SYNCHRONIZATION SETTINGS

To make Synchronization settings:

- Click the **Administration** tab.
- Click the **Background Activities** icon.

The list of background activities appears.

- Click the **Settings** button.
- Click the Synchronization Rate drop-down menu and choose a rate:
 - Low** – Fewer system resources to Synchronization, more to data read/write operations.
 - Medium** – Balances system resources between Synchronization and data read/write operations.
 - High** – More system resources to Synchronization, fewer to data read/write operations.
- Click the **Confirm** button.

BATTERY RECONDITIONING

Batteries maintain power to the controller cache in the event of a power failure, thus protecting any data that has not been written to a physical drive. Reconditioning is the action of discharging and recharging a battery to preserve its capacity and performance.

MANAGING STORAGE SERVICES

Storage service management includes:

- "Viewing a List of Services" on page 78
- "Email Service" on page 78
- "Webserver Service" on page 79
- "Telnet Service" on page 80
- "SSH Service" on page 80
- "SNMP Service" on page 81
- "Netsend Service" on page 82

VIEWING A LIST OF SERVICES

This feature displays all software services running on the RAID subsystem. See the table below.

To view the list of software services:

1. Click the **Administration** tab.
2. Click the Services icon.

The Services list displays the Status and Start Type of the services available. These services are described in the sections that follow.

EMAIL SERVICE

Email service enables the RAID subsystem to send you Email messages about events and status changes. By default, Email service is set to Automatic.

STOPPING EMAIL SERVICE

To stop the Email service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Stop** button.
4. Click the **Confirm** button.

To start the Email service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Start** button.

RESTARTING EMAIL SERVICE

To restart the Email service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Restart** button.

MAKING EMAIL SETTINGS

To change Email service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Settings** button.
4. Make settings changes as required:
 - Choose a startup type,
 - Automatic – (default) Starts and runs with the subsystem.
 - Manual – You start the service when you need it.
 - SMTP Server IP address.
 - SMTP Authentication – The Yes option enables authentication. The No option disables.
 - SMTP Authentication Username – Required if SMTP authentication is enabled.
 - SMTP Authentication Password – Required if SMTP authentication is enabled.
 - Email Sender (From) Address – The sender’s name shown on notification messages.
 - Email Subject – The subject line of the notification message.
5. Click the **Save** button.
6. Click the **Confirm** button.



Note

To verify your settings, send a test message.

SLP SERVICE

Service Location Protocol (SLP) discovers services over the Internet. SLP applies to IPv4 protocol only.

STOPPING SLP SERVICE

To stop the SLP service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Stop** button.
4. Click the **Confirm** button.

To start the SLP service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Start** button.

RESTARTING SLP SERVICE

To restart the SLP service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Restart** button.

MAKING SLP SETTINGS

To change SLP service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Settings** button.
4. Choose a startup type:
 - Automatic – (default) Starts and runs with the subsystem.
 - Manual – You start the service when you need it.
5. Click the **Save** button.
6. Click the **Confirm** button.

WEBSERVER SERVICE

Webserver service connects the WebPAM PROe interface to the RAID subsystem though your browser.

STOPPING WEBSERVER SERVICE

To stop the Webserver service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Webserver service and click the **Stop** button.
4. Click the **Confirm** button.

To start the Webserver service after stopping it:

1. Click the **Administration** tab.
2. Click the Services icon.
3. Click the Webserver service and click the **Start** button.

RESTARTING WEBSERVER SERVICE

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Webserver service and click the **Restart** button.

MAKING WEBSERVER SETTINGS

To change Webserver service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Webserver service and click the **Settings** button.
4. Make settings changes as required:
 - Choose a startup type,
 - Automatic – (default) Starts and runs with the subsystem.
 - Manual – You start the service when you need it.
 - Session Time Out – Default is 24 minutes.
5. Click the **Save** button.
6. Click the **Confirm** button.

TELNET SERVICE

Telnet service enables you to access the RAID subsystem’s Command Line Interface (CLI) through a network connection.

STOPPING TELNET SERVICE

To stop the Telnet service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **Telnet** service and click the **Stop** button.
4. Click the **Confirm** button.

To start the Telnet service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **Telnet** service and click the **Start** button.

RESTARTING TELNET SERVICE

To restart the Telnet service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click Telnet service and click the **Restart** button.

MAKING TELNET SETTINGS

To change Telnet service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **Telnet** service and click the **Settings** button.
4. Make settings changes as required:
 - Choose a startup type,
Automatic – Starts and runs with the subsystem.
Manual – (default) You start the service when you need it.
 - Port number – Default is 2300.
 - Max Number of Concurrent Connections – Default is 4.
Maximum number is 4.
 - Session Time Out – Default is 24 minutes.

5. Click the **Save** button.
6. Click the **Confirm** button.

SSH SERVICE

Secure Shell (SSH) service enables you to access the subsystem’s Command Line Interface (CLI) through a network connection.

STOPPING SSH SERVICE

To stop SSH service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SSH** service and click the **Stop** button.
4. Click the **Confirm** button.

To start SSH service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click SSH service and click the **Start** button.

RESTARTING SSH SERVICE

To restart SSH service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SSH** service and click the **Restart** button.

MAKING *SSH* SETTINGS

To change SSH service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SSH service and click the **Settings** button.
4. Make settings changes as required:
 - Choose a startup type,
Automatic – (default) Starts and runs with the subsystem.
Manual – You start the service when you need it.
 - Port number - Default is 22.
 - Max Number of Concurrent Connections – Default is 4.
Maximum number is 4.
 - Session Time Out - Default is 24 minutes.
5. Click the **Save** button.
6. Click the **Confirm** button.

SNMP SERVICE

Simple Network Management Protocol (SNMP) service enables the SNMP browser to obtain information from the RAID subsystem. The Trap Sink is where SNMP events are sent and can be viewed.

STOPPING *SNMP* SERVICE

To stop the SNMP service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Stop** button.
4. Click the **Confirm** button.

To start the SNMP service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Start** button.

RESTARTING *SNMP* SERVICE

To restart the SNMP service:

1. Click the Administration tab.
2. Click the Services icon.
3. Click the SNMP service and click the Restart button.

MAKING *SNMP* SETTINGS

To change SNMP service settings:

1. Click the Administration tab.
2. Click the Services icon.
3. Click the SNMP service and click the Settings button.
4. Make settings changes as required:
 - Choose a startup type,
Automatic – (default) Starts and runs with the subsystem.
Manual – You start the service when you need it.
 - Port Number – Default is 161.
 - System Name – No default.
 - System Location – Default is USA.
 - System Contact – Default is admin@yourcompany.com.
 - Read Community – Default is public.
 - Write Community – Default is private. No changes are possible.
5. Click the Save button.
6. Click the **Confirm** button.

ADDING AN *SNMP* TRAP SINK

To add a trap sink:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Settings** button.
4. **Enter** a trap sink server IP address in the field provided.
5. Choose a trap filter (event severity level).

See the table on the next page.
6. Click the **Add** button.
7. Click the **Confirm** button.

DELETING AN **SNMP TRAP SINK**

To delete a trap sink:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Settings** button.
4. In the **Trap Sink** list and click the Trap Sink you want to delete.
5. Click the **Trash** icon.

The trap sink is deleted.

6. Click the **Save** button.
7. Click the **Confirm** button.

NETSEND SERVICE

Netsend service sends RAID subsystem events in the form of text messages to the Host PC and other networked PCs configured to receive Netsend event messages by setting up Netsend server accounts.

This service is set to Manual startup by default. It does not run unless you start it manually or change the startup type to Automatic.

STARTING **NETSEND SERVICE**

To restart the Netsend service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the Start button.

STOPPING **NETSEND**

To stop the Netsend service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Netsend service and click the Stop button.
4. Click the **Confirm** button.

RESTARTING **NETSEND SERVICE**

To start the Netsend service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Start** button.

MAKING NETSEND SETTINGS

To change Netsend service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Settings** button.
4. Choose a startup type,
 - Automatic – Starts and runs with the subsystem.
 - Manual – (default) You start the service when you need it.
5. Click the **Save** button.
6. Click the **Confirm** button.

ADDING NETSEND SERVER ACCOUNTS

To add a Netsend server account:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Settings** button.
4. **Enter** the recipient server IP address in the field provided.
5. Choose a recipient filter (event severity level).

See the table on the next page.

6. Click the **Add** button.

The recipient server is added to the list.

7. Click the **Save** button.
8. Click the **Confirm** button.

Event Severity Levels	
Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

DELETING NETSEND SERVER ACCOUNTS

To delete a Netsend server account:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Settings** button.
4. In the **Message Event Severity Filter** list, click the recipient server you want to delete.
5. Click the **Trash** icon.

The recipient server is deleted.

6. Click the **Save** button.
7. Click the **Confirm** button.

WORKING WITH THE EVENT VIEWER

Working with the Event Viewer includes the following functions:

- "Viewing Runtime Events" on page 84
- "Saving Runtime Events" on page 84
- "Clearing Runtime Events" on page 84
- "Viewing NVRAM Events" on page 85
- "Saving NVRAM Events" on page 85
- "Clearing NVRAM Events" on page 85

The Event Viewer displays log of subsystem events. Events are classified as:

- Runtime Events – A list of and information about the 1023 most recent runtime events recorded since the subsystem was started.
- **NVRAM Events** – A list of and information about the most important events over multiple subsystem startups. NVRAM events are stored in non-volatile memory.

Event Severity Levels	
Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

VIEWING RUNTIME EVENTS

To display Runtime Events:

1. Click the **Administration** tab.
2. Click the **Events** icon.

The log of Runtime Events appears. Events are added to the top of the list. Each item includes:

- **Index number** – Begins with 0 at system startup.
- **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
- **Event ID** – Hexadecimal code for the specific event
- **Severity** – See the Table above.
- **Timestamp** – Date and time the event happened.
- **Description** – A description of the event in plain language.

3. Press the up and down arrow keys to scroll through the log.

SAVING RUNTIME EVENTS

This feature saves a plain text file of runtime events to your host PC or server using your browser.

To save the Runtime Events log:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **Save** button.
4. Follow your browser’s procedure to save the event file to the desired location.

CLEARING RUNTIME EVENTS

To clear the Runtime Events log:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **Clear** button.
4. In the Confirmation box, type the word “**confirm**” in the field provided and click the Confirm button.

VIEWING NVRAM EVENTS

This screen displays a list of and information about the most important events over multiple subsystem startups.

To display NVRAM events:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **NVRAM Events** button.

The log of NVRAM Events appears. Events are added to the top of the list. Each item includes:

- **Index number** – Begins with 0 at system startup.
- **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
- **Event ID** – Hexadecimal code for the specific event
- **Severity** – See the Table on the previous page.
- **Timestamp** – Date and time the event happened.
- **Description** – A description of the event in plain language.

4. Press the up and down arrow keys to scroll through the log.

SAVING NVRAM EVENTS

This feature saves a plain text file of NVRAM events to your host PC or server using your browser.

To save NVRAM Events:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **NVRAM Events** button.
4. Click the **Save** button.
5. Follow your browser’s procedure to save the event file to the desired location.

CLEARING NVRAM EVENTS

To clear the Runtime Events log:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **Clear** button.
4. In the Confirmation box, type the word “confirm” in the field provided and click the **Confirm** button.

MONITORING PERFORMANCE

Performance monitoring includes:

- Monitoring I/O Performance (below)
- Monitoring PSU Wattage (page 86)

MONITORING I/O PERFORMANCE

The Performance Monitor displays real-time performance statistics for logical drives, physical drives, and Fibre Channel or iSCSI data ports. The vertical scale adjusts dynamically to accommodate the statistical data.

Because it reports performance in real-time, to see data in the monitor, there must be I/O data activity taking place between the Vess R2600 subsystem and the Host.

To monitor performance:

1. Click the **Administration** tab.
2. Click the **Performance Monitor** icon.
3. Under **Logical Drive**, choose the metric you want to see from the **Measurement** drop-down menu.
 - Bandwidth in MB/s
 - Cache usage by %
 - Dirty cache usage by %
 - Maximum latency in ms
 - Average latency in ms
 - Minimum latency in ms
 - I/Os per second
4. Click the **Select Logical Drives** button and check the boxes for the logical drives you want to see.
 - Total of all logical drives
 - Up to 8 individual logical drives
5. Under **Physical Drive**, choose the metric you want to see from the **Measurement** drop-down menu.
 - Bandwidth in MB/s
 - Maximum latency in ms
 - Average latency in ms
 - Minimum latency in ms
 - I/Os per second

6. Click the Select **Physical Drives** button and check the boxes for the physical drives you want to see.
- Total of all physical drives
 - Up to 8 individual physical drives
7. Under **Port**, choose the metric you want to see from the **Measurement** drop-down menu.
- Bandwidth in MB/s
 - Maximum latency in ms
 - Average latency in ms
 - Minimum latency in ms
 - I/Os per second
8. Click the **Select Ports** button and check the boxes for the ports you want to see:
- Total of all ports
 - Up to 8 individual ports

Since the **Performance Monitor** is a real-time display, it does not accumulate information and there is no clear or save function.

To save performance statistics for analysis or troubleshooting, save a **Service Report**, open the report, and look under **Statistic Info**.

MONITORING PSU WATTAGE

The PSU Wattage Monitor displays real-time performance statistics for logical drives, the input power of all enclosures and the input power of an individual. The vertical scale adjusts dynamically to accommodate the statistical data.

Because it reports performance in real-time, to see data in the monitor, there must be I/O data activity taking place between the Vess R2600 subsystem and the Host.

To monitor performance and power use:

1. Click the **Administration** tab.
2. Click the **PSU Wattage Monitor** icon.
3. Under **Input Power of an individual Enclosure**, click the **Select Enclosures** button and check the boxes for the enclosures you want to see.

Since the PSU Wattage Monitor is a real-time display, it does not accumulate information and there is no clear or save function.

To save performance and power statistics for analysis or troubleshooting, save a Service Report, open the report, and look under PSU Wattage Info.

MANAGING PHYSICAL DRIVES

Physical drive management includes:

For physical disk troubleshooting, see "Physical Drive Problems" on page 213.

- Viewing a List of Physical Drives (below)
- Viewing Physical Drive Information (below)
- "Making Global Physical Drive Settings" on page 88
- "Making Individual Physical Drive Settings" on page 88
- "Viewing Physical Drive Statistics" on page 88
- "Viewing Physical Drive SMART Log Information" on page 89
- "Saving the Physical Drive SMART Log" on page 89
- "Locating a Physical Drive" on page 89
- "Forcing a Physical Drive Offline" on page 90
- "Clearing a Stale or a PFA Condition" on page 90
- "Updating Firmware on a Physical Drive" on page 90




VIEWING A LIST OF PHYSICAL DRIVES

To view a list of physical drives in the RAID system:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.

The list of enclosures and the physical drives inside them displays.

Physical drive information includes:

- ID – ID number of the physical drive
- Status – Green check , yellow ! , and red X  icons
- Model – Make and model of the drive
- Type – SAS or SATA, HDD or SSD
- Location – Enclosure number and slot number
- Configuration – Array number and sequence number, spare number, unconfigured, or stale configuration
- Capacity – In GB

VIEWING PHYSICAL DRIVE INFORMATION

To view physical drive information:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button.

Physical drive information includes:

- | | |
|--|--|
| <ul style="list-style-type: none">• Model - Model of PROMISE system• Physical Drive ID – ID number of the physical drive• Location – Enclosure number and slot number• Alias – If assigned• Physical Capacity – Total capacity in GB• Configurable Capacity – Usable capacity in GB• Used Capacity – Capacity actually used in GB• Block Size – Typically 512 Bytes• Operational Status – OK is normal, Stale, PFA, Dead | <ul style="list-style-type: none">• Configuration – Array number and sequence number, spare number, Model – Make and model of the drive• Drive Interface – SATA 1.5Gb/s or 3Gb/s, SAS 3Gb/s or 6Gb/s• Serial Number – Serial number of the drive• Firmware Version – Firmware version on the drive• Protocol Version – ATA/ATAPI or SCSI protocol version• Visible To – Controllers that can access this physical drive |
|--|--|

Advanced information for SATA physical drives includes:

- | | |
|---|--|
| <ul style="list-style-type: none">• Write Cache – Enabled or disabled• Read Look Ahead Cache – Enabled or disabled• Read Cache Support – Yes or No• SMART Feature Set – Yes or No• SMART Self Test – Yes or No• SMART Error Logging – Yes or No• Command Queuing Support – TCQ or NCQ• Command Queuing – Enabled or disabled | <ul style="list-style-type: none">• Queue Depth - Number of commands• Maximum Multiple DMA Mode Supported• Maximum Ultra DMA Mode Supported• DMA Mode• Power Saving Level – Enabled or disabled• APM Support – Standby or Active• Medium Error Threshold• Drive Temperature• Drive Reference Temperature |
|---|--|

Advanced information for SAS physical drives includes:

- Read Cache – Enabled or disabled
- Read Cache Support – Yes or No
- Write Cache – Enabled or disabled
- Write Cache Support – Yes or No
- Enable Read Look Ahead Support – Yes or No
- Read Look Ahead Cache – Enabled or disabled
- Command Queuing – Enabled or disabled
- Command Queuing Support – Yes or No
- WWN – World Wide Name
- Port 1 Negotiated Physical Drive Speed
- Port 1 SAS Address
- Port 2 Negotiated Physical Drive Speed
- Port 2 SAS Address
- Drive Temperature in °C
- Drive Reference Temperature in °C
- Power Saving Level – Enabled or disabled
- Medium Error Threshold
- SAS SATA Bridge Firmware Version
- SAS SATA Bridge Boot Loader Version

MAKING GLOBAL PHYSICAL DRIVE SETTINGS

To make global physical drive settings:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the **Global Physical Drive Settings** button.
4. Check the boxes to enable, uncheck to disable.

For SATA drives:

- Enable Write Cache
- Enable Read Look Ahead Cache
- Enable Command Queuing

For SAS drives:

- Enable Write Cache
- Enable Read Look Ahead Cache
- Enable Read Cache

5. Click the **Save** button.

MAKING INDIVIDUAL PHYSICAL DRIVE SETTINGS

To make individual physical drive settings:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **Settings** button.
4. On the **Settings** tab:
 - **Enter**, change, or delete the alias in the **Alias** field.
5. On the **SMART Log Settings** tab:
 - Check the box to enable the SMART log.
6. Click the **Save** button.

VIEWING PHYSICAL DRIVE STATISTICS

To view physical drive statistics:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button.
4. Click the **Statistics** tab.

Physical drive statistics include

- Data Transferred
- Read Data Transferred
- Write Data Transferred
- Errors - Number of errors
- Non Read/Write Errors
- Read Errors
- Write Errors
- I/O Request – Number of requests
- Non Read/Write Request – Number of requests
- Read I/O Request – Number of requests
- Write I/O Request – Number of requests
- Statistics Start Time – Time and date
- Statistics Collection Time – Time and date
- Avg Response Time Ctrl 1 – Controller 1 average response time
- Avg Response Time Ctrl 2 – Controller 2 average response time
- Max Response Time Ctrl 1 – Controller 1 maximum response time
- Max Response Time Ctrl 2 – Controller 2 maximum response time

To clear physical drive statistics, see "Clearing Statistics" on page 55.

VIEWING PHYSICAL DRIVE SMART LOG INFORMATION

To view physical drive SMART Log information:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button.
4. Click the **SMART Log** tab.

SMART Log information includes:

- In progress
- SMART Support – Yes or no, depends on the drive
- SMART Log Enabled – Enabled or disabled, see Note below
- SMART Health status – OK is normal
- SCT Status Version
- SCT Version
- SCT Support Level
- Device State
- Current Temperature
- Power Cycle Min Temperature
- Power Cycle Max Temperature
- Lifetime Min Temperature
- Lifetime Max Temperature
- Under Temperature Limit Count
- Over Temperature Limit Count

If the SMART Log is disabled, see Making Controller Settings.

SAVING THE PHYSICAL DRIVE SMART LOG

To save the physical drive SMART Log:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button.
4. Click the **SMART Log** tab.
5. Click the **Save Advanced SMART Log** button.

Your browser saves a text file containing the SMART Log to its designated download folder.

LOCATING A PHYSICAL DRIVE

This feature causes the drive carrier LEDs to blink for one minute to assist you in locating the physical drive, and is supported by RAID subsystems and JBOD expansion units.

To locate a physical drive:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **Locate** button.

The drive carrier status LED flashes for one minute.

FORCING A PHYSICAL DRIVE OFFLINE

This feature applies only to physical drives assigned to disk arrays.



Caution

Forcing a physical drive offline is likely to cause data loss. Back up your data before you proceed. Use this function only when required.



Important

Forcing a physical drive offline causes your logical drives to become degraded. If Auto Rebuild is enabled and a spare drive is available, the disk array begins rebuilding itself automatically.

To force a physical drive offline:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the **down arrow** button to list the physical drives in the enclosure.
4. Mouse over the physical drive you want to force offline.
5. Click the **Force Offline** button.
6. In the **Confirmation** box, type the word “**confirm**” in the field provided and click the **Confirm** button.

CLEARING A STALE OR A PFA CONDITION

This procedure is used to clear configuration data on a physical drive; or if the physical drive is stale or the has errors putting it in PFA status.

Stale – The physical drive contains obsolete disk array information.

PFA – The physical drive has errors resulting in a prediction of failure.

Be sure you have corrected the condition by a physical drive replacement, rebuild operation, etc., first. Then clear the condition.

To clear a **Stale** or a **PFA** condition:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **Clear** button.

If the physical drive has both a Stale condition and a PFA condition, the first click removes the Stale condition. Click the **Clear** button a second time to remove the PFA condition.

UPDATING FIRMWARE ON A PHYSICAL DRIVE

This feature applies only to PROMISE-supported physical drives. For a list of supported drives, go to <http://www.promise.com/support/>.

Then see "Updating Physical Drive Firmware" on page 175.

If you have physical drives in your RAID system that are not PROMISE-supported, follow the firmware update procedure from the drive manufacturer.

MANAGING DISK ARRAYS

For disk array troubleshooting, see "Disk Array and Logical Drive Problems" on page 213.

Disk array management includes:

- "Viewing a List of Disk Arrays" on page 91
- "Creating a Disk Array Manually" on page 92
- "Creating a Disk Array with the Wizard" on page 92
- "Deleting a Disk Array" on page 93
- "Making Disk Array Settings" on page 93
- "Running PDM on a Disk Array" on page 94
- "Preparing a Disk Array for Transport" on page 95
- "Rebuilding a Disk Array" on page 95


VIEWING A LIST OF DISK ARRAYS

To view a list of disk arrays:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

The list of disk arrays appears.

Disk array information includes:

- **ID** – DA0, DA1, DA2, etc.
- **Alias** – If assigned
- **Status** – A green check  icon means OK
- **Capacity** – Data capacity of the array
- **Free Capacity** – Unconfigured or unused capacity on the physical drives
- **Media Patrol** – Enabled or disabled on this array
- **No. of Logical Drives** – The number of logical drives on this array

VIEWING DISK ARRAY INFORMATION

To view disk array information:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

The list of disk arrays appears.

3. Click the disk array you want, then click the **View** button.

Array information displays, including:

- **ID** – DA0, DA1, DA2, etc.
- **Alias** – If assigned
- **Operational Status** – OK is normal
- **Media Patrol** – Enabled or disabled on this array
- **PDM** – Enabled or disabled on this array
- **Total Physical Capacity** – Data capacity of the array
- **Configurable Capacity** – Maximum usable capacity of the array
- **Free Capacity** –Unconfigured or unused capacity on the physical drives
- **Max Contiguous Free Capacity - The largest contiguous free capacity available.**
- **Current Power Saving Level** - Default is disabled
- **Number of Physical Drives** – The number of physical drives in this array
- **Number of Logical Drives** – The number of logical drives on this array
- **Max Contiguous Free Capacity** – Unconfigured or unused capacity in contiguous sectors on the physical drives
- **Available RAID Levels** – RAID levels you can specify on this array

DISK ARRAY OPERATIONAL STATUS

- **OK** – This is the normal state of a logical drive. When a logical drive is Functional, it is ready for immediate use. For RAID Levels other than RAID 0 (Striping), the logical drive has full redundancy.
- **Synchronizing** – This condition is temporary. Synchronizing is a maintenance function that verifies the integrity of data and redundancy in the logical drive. When a logical drive is Synchronizing, it functions and your data is available. However, access is slower due to the synchronizing operation.
- **Critical/Degraded** – This condition arises as the result of a physical drive failure. A degraded logical drive still functions and your data is still available. However, the logical drive has lost redundancy (fault tolerance). You must determine the cause of the problem and correct it.
- **Rebuilding** – This condition is temporary. When a physical drive has been replaced, the logical drive automatically begins rebuilding in order to restore redundancy (fault tolerance). When a logical drive is rebuilding, it functions and your data is available. However, access is slower due to the rebuilding operation.
- **Transport Ready** – After you perform a successful Prepare for Transport operation, this condition means you can remove the physical drives of this disk array and move them to another enclosure or different drive slots. After you relocate the physical drives, the disk array status shows OK.

CREATING A DISK ARRAY MANUALLY

This feature creates a disk array only. You can also use the Wizard to create a disk array with logical drives and spare drives at the same time.

This action requires **Super User** or **Power User** privileges.

To create a disk array:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.
3. Click the **Create Disk Array** button.
4. Accept the defaults or make changes:
 - **Enter** an alias in the **Alias** field
Maximum of 32 characters; letters, numbers, space between characters, and underline.
 - **Media Patrol** – Uncheck to disable on this array.
 - **PDM** – Uncheck to disable on this array.
 - **Power Management** – Uncheck to disable on this array.
 - **Choose a media type** – Hard disk drive (HDD) or solid state drive (SSD)



Important

All physical drives in an array must be the same media type, i.e. all HDD or all SSD.

5. In the **Select Physical Drives** diagram, click the drives to add them to your array. Look for drives with a green LED dark, a blue LED lit, and no crosshatching over the carrier.

The ID numbers of the chosen drives appear in the field below the diagram.

6. When you have finished your settings and choices, click the **Submit** button.

The new array appears in the list.

If you are done creating disk arrays, click the **Finish** button.

To create additional disk arrays, click the **Create More** button.

After you create a disk array, create a logical drive on it. See "Creating a Logical Drive Manually" on page 97.

CREATING A DISK ARRAY WITH THE WIZARD

The Wizard creates disk arrays and logical drives automatically. It has four options.

- **Automatic** – Creates a new disk array following a default set of parameters. Creates a hot spare drive for all RAID levels except RAID 0, when five or more unconfigured physical drives are available. You can accept or reject the proposed arrangement but you cannot modify it. See instructions in "Wizard - Automatic Configuration" on page 46.
- **Advanced** – Enables you to specify all parameters for a new disk array, logical drives and spare drives. See instructions in "Wizard - Advanced Configuration" on page 47.

DELETING A DISK ARRAY



Caution

If you delete a disk array, you also delete any logical drives that belong to it, along with the data in those logical drives. Back up any important data before deleting a disk array.

This action requires **Administrator** or **Super User** privileges.

To delete a disk array:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.
3. Click the disk array you want, then click the **Delete** button.
4. In the **Confirmation** box, type the word “**confirm**” in the field provided and click the **Confirm** button.

LOCATING A DISK ARRAY

This feature causes the drive carrier LEDs to flash for one minute to assist you in locating the physical drives that make up this disk array.

To locate a disk array:


1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

The list of disk arrays appears.

3. Click the disk array you want, then click the **Locate** button.

The drive carrier status LEDs for the array flash for one minute.

Figure 4 - 6: Drive carrier status LED



MAKING DISK ARRAY SETTINGS

To make disk array settings:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

The list of disk arrays appears.

3. Click the disk array you want, then click the **Settings** button.
4. Make settings changes as required:
 - **Enter**, change or delete the alias in the **Alias** field
Maximum of 32 characters; letters, numbers, space between characters, and underline.
 - **Media Patrol** – Check to enable, uncheck to disable on this array.
 - **PDM** – Check to enable, uncheck to disable on this array.
 - **Power Management** – Check to enable, uncheck to disable on this array.
5. Click the **Save** button.



Notes

You can also enable or disable Media Patrol for the entire RAID system. See Making Media Patrol Settings.

HDD Power Saving must be enabled on the RAID controller for the Power Management settings to be effective. See Making Controller Settings.

Power Management functions are limited to the features your HDDs actually support.

RUNNING MEDIA PATROL ON A DISK ARRAY

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. If Media Patrol encounters a critical error, it triggers PDM if PDM is enabled on the disk array.

For more information, see "Media Patrol" on page 75. Also see"PDM" on page 76.

RUNNING MEDIA PATROL

To run **Media Patrol**:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over Media Patrol and click the **Start** button.

STOPPING, PAUSING OR RESUMING MEDIA PATROL

To stop, pause or resume **Media Patrol**:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Mouse-over Media Patrol and click the **Stop, Pause** or **Resume** button.

RUNNING PDM ON A DISK ARRAY

Predictive Data Migration (PDM) is the migration of data from the suspect disk drive to a spare disk drive.

For more information, see "PDM" on page 76.

RUNNING PDM

To run PDM on a disk array:

1. Click the **Administration** tab.
2. Click the Background **Activities** icon.

The list of background activities appears.

3. Mouse-over PDM and click the **Start** button.
4. From the **Source Physical Drive** drop-down menu, choose a Source disk array and physical drive.
5. From the **Target Physical Drive** drop-down menu, choose a Target physical drive.
6. Click the **Confirm** button.

STOPPING, PAUSING OR RESUMING PDM

To stop, pause or resume PDM:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over PDM and click the **Stop, Pause**, or **Resume** button.

You can also enable or disable PDM on individual disk arrays. See "Making Disk Array Settings" on page 93.

PREPARING A DISK ARRAY FOR TRANSPORT

This feature prepares the physical drives that make up the disk array to be removed from the enclosure and installed in a different location.

To prepare a disk array for transport:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

The list of disk arrays appears.

3. Click the disk array you want, then click the **Transport** button.
4. Click the **Confirm** button.

The status changes to **Transport Ready**.

5. Remove the physical drives and install them in their new location.

For more information, see “Installing Physical Drives” on page 16.

REBUILDING A DISK ARRAY

When you rebuild a disk array, you are actually rebuilding the data on one physical drive.

If there is no spare drive of adequate capacity and the **Auto Rebuild** function is DISABLED, you must replace the failed drive with an unconfigured physical drive, then perform a Manual Rebuild. See "Making Rebuild Settings" on page 76.



Important

If your replacement disk drive was formerly part of a different disk array or logical drive, you must clear the configuration data on the replacement drive before you use it. See "Clearing a Stale or a PFA Condition" on page 90.

PERFORMING A MANUAL REBUILD

To perform a manual rebuild:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.



The list of background activities appears.

3. Mouse-over Rebuild and click the **Start** button.
4. From the Source Physical Drive drop-down menu, choose a Source disk array and physical drive.

Arrays have an ID No. Physical drives have a Seq. No.(sequence number)

5. From the Target Physical Drive drop-down menu, choose a Target physical drive.
6. Click the **Confirm** button.

When the disk array is rebuilding:

- The disk array shows a green check  icon and **Rebuilding** status.
- Logical drives under the disk array continue to show a yellow !  icon and **Critical** status.

STOPPING, PAUSING OR RESUMING A REBUILD

To stop, pause or resume a Rebuild:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background appears.

3. Mouse-over Rebuild and click the **Stop, Pause, or Resume** button.

MANAGING LOGICAL DRIVES

Logical drive management includes:

- "Viewing a List of Logical Drives" on page 96
- "Viewing Logical Drive Information" on page 96
- "Viewing Logical Drive Statistics" on page 97
- "Viewing Logical Drive Check Tables" on page 97
- "Creating a Logical Drive Manually" on page 97
- "Deleting a Logical Drive" on page 98
- "Making Logical Drive Settings" on page 98
- "Locating a Logical Drive" on page 98
- "Initializing a Logical Drive" on page 99
- "Redundancy Check on a Logical Drive" on page 99
- "Migrating a Logical Drive’s RAID Level" on page 100
- "Creating a LUN Clone" on page 100


VIEWING A LIST OF LOGICAL DRIVES

To view a list of logical drives:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

The list of logical drives appears.

Logical Drive information includes:

- **ID** – LD0, LD1, LD2, etc.
- **Alias** – If assigned.
- **Status** – A green check
-  icon means OK.
- **Capacity** – Data capacity of the logical drive.
- **RAID Level** – Set when the logical drive was created.
- **Stripe** – Set when the logical drive was created.
- **Cache Policy** – Read cache and Write cache settings.
- **Array ID** – ID number of the disk array where this logical drive was created.

VIEWING LOGICAL DRIVE INFORMATION

To view logical drive information:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

The list of logical drives appears.

3. Click the logical drive you want, then click the **View** button.

Logical Drive information displays, including:

- **ID** – LD0, LD1, LD2, etc.
- **Alias** – If assigned
- **Array ID** — ID number of the disk array where this logical drive was created
- **RAID Level** – Set when the logical drive was created
- **Operational Status** – OK means normal
- **Capacity** – Data capacity of the logical drive
- **Number of Axles** – For RAID 10, 2 axles. For RAID 50 and 60, 2 or more axles
- **Physical Capacity** – Data capacity of the physical drives
- **Number of Physical Drives** – The number of physical drives in the disk array
- **Stripe size** – Set at logical drive creation
- **Read Policy** – Adjustable
- **Sector size** – Set at logical drive creation
- **Write Policy** – Adjustable
- **Preferred Controller ID** – For RAID subsystems with dual controllers
- **Tolerable Number of Dead Drives** – Number of physical drives that can fail without the logical drive going offline
- **Host Accessibility** - Normal, read-only, write-only, or not visible to host
- **Synchronized** – A new logical drive shows “No” until synchronizing is completed. See Synchronization
- **Parity Pace** – Pertains to some RAID levels
- **WWN** – Worldwide Name, a unique identifier assigned to this logical drive
- **Codec Scheme** – Pertains to some RAID levels
- **Serial Number** – Assigned to this logical drive
- **ALUA Access State for Ctrl 1** - Active, optimized or standby
- **ALUA Access State for Ctrl 2** - Active, optimized or standby
- **Association State** - for LUN clone
- **Storage Service Status** - for LUN clone
- **PerfectRebuild** - Enable or disable

VIEWING LOGICAL DRIVE STATISTICS

To view logical drive statistics:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

The list of logical drives appears.

3. Click the logical drive you want, then click the **View** button.
4. Click the **Statistics** tab.

Logical Drive statistics display, including:

- **Data Transferred** – In bytes
- **Read Data Transferred** – In bytes
- **Write Data Transferred** – In bytes
- **Errors**
- **Read Errors**
- **Write Errors**
- **I/O Requests**
- **Non-Read/Write I/O Requests**
- **Read I/O Requests**
- **Write I/O Requests**
- **Statistics Start Time**
- **Statistics Collection Time**

To clear physical drive statistics, see "Clearing Statistics" on page 55.

VIEWING LOGICAL DRIVE CHECK TABLES

This feature enables you to view error tables. Use this information to evaluate the integrity of the logical drive and to determine whether corrective action is needed.

To view logical drive check tables:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

The list of logical drives appears.

3. Click the logical drive you want, then click the **Check Table** button.

4. Choose an option:
 - **All** – All errors. The default choice.
 - **Read Check** – Read errors for this logical drive.
 - **Write Check** – Write errors for this logical drive.
 - **Inconsistent Block** – Inconsistent blocks for this logical drive. Mirror data for RAID Levels 1, 1E and 10 or Parity data for RAID Levels 3, 5, 6, 30, 50, and 60. Identified by the Redundancy Check.

The Check Table lists:

- **Entry Number** – A number assigned to each block of entry.
- **Table Type** – Read Check, Write Check or Inconsistent Block.
- **Start Logical Block Address** – LBA of the first block for this entry.
- **Count** – Number of errors or continuous blocks starting from this LBA.

To clear the check tables, see "Clearing Statistics" on page 55.

CREATING A LOGICAL DRIVE MANUALLY

This feature creates a logical drive only. You can also use the Wizard to create a disk array with logical drives and spare drives at the same time. **See Creating a Disk Array with the Wizard.**

This action requires **Super User** or **Power User** privileges.

To create a logical drive manually:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.
3. Click the **Create Logical Drive** button.
4. Click the option button of the disk array you want to use and click the Next button.
5. Optional. **Enter** an alias in the **Alias** field.

Maximum of 32 characters; letters, numbers, space between characters, and underline.

6. Choose a RAID level.

The choice of RAID levels depends the number of physical drives in the disk array.

7. RAID 30, 50 and 60 only. Specify the number of axles for your array.
8. In the **Capacity** field, accept the default maximum capacity or enter a lesser capacity and size in MB, GB or TB.

Any remaining capacity is available for an additional logical drive.

9. For each of the following items, accept the default or change the settings as required:
- Choose a Stripe size.
64 KB, 128 KB, 256 KB, 512 KB, and 1 MB are available.
 - Choose a Sector size.
512 B, 1 KB, 2 KB, and 4 KB are available.
 - Choose a Read (cache) Policy.
Read Cache, Read Ahead, and No Cache are available.
 - Choose a Write (cache) Policy.
Write Back and Write Through (Thru) are available.
 - Check box of Perfect Rebuild Enable / Disable Perfect Rebuild

10. Click the **Add** button.

The new logical drive appears on the list at the right.

If there is capacity remaining, you can create an additional logical drive.

11. When you are finished, click the **Submit** button.

The new logical drive or drives appear in the logical drive list.

New logical drives are automatically synchronized. You can access the logical drive during synchronization.

DELETING A LOGICAL DRIVE



Caution

If you delete a logical drive, you also delete all the data in the logical drive. Back up any important data before deleting the logical drive.

This action requires **Administrator** or **Super User** privileges.

To delete a logical drive:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.
3. Click the logical drive you want, then click the **Delete** button.
4. In the Confirmation box, type the word “**confirm**” in the field provided and click the **Confirm** button.

MAKING LOGICAL DRIVE SETTINGS

To make logical drive settings:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

The list of logical drives appears.

3. Click the logical drive you want, then click the **Settings** button.
4. Make settings changes as required:
 - **Enter**, change, or delete the alias in the Alias field.
Maximum of 32 characters; letters, numbers, space between characters, and underline.
 - Choose a Read (cache) Policy.
Read Cache, Read Ahead, and No Cache are available.
 - Choose a Write (cache) Policy.
Write Back and Write Through (Thru) are available.
 - Check box of Perfect Rebuild Enable / Disable Perfect Rebuild Note that if Perfect Rebuild is disabled, it cannot be enabled again on the LD.
5. Click the **Save** button.

For more information, see "Cache Policy" on page 195.



Note

The Write Cache is always set to WriteThru when Read Cache is set to NoCache.

LOCATING A LOGICAL DRIVE

This feature causes the drive carrier LEDs to flash for one minute to assist you in locating the physical drives that make up this logical drive.

To locate a logical drive:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

The list of logical drives appears.

3. Click the logical drive you want, then click the **Locate** button.

The drive carrier status LEDs of the disk carriers making up the logical drive flash for one minute.

Figure 4 - 7: Drive carrier status LED



INITIALIZING A LOGICAL DRIVE

Initialization is normally done to logical drives after they are created from a disk array.



Warning

When you initialize a logical drive, all the data on the logical drive is lost. Backup any important data before you initialize a logical drive.

To initialize a logical drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over Initialization and click the **Start** button.
4. Check the box to the left of the logical drive you want to initialize.
5. Choose the initialization option you want:
 - **Quick Initialization** – Check the box and enter a value in the Quick Initialization Size field. This value is the size of the initialization blocks in MB.
 - **Full Initialization** – Do not check the box. **Enter** a hexadecimal value in the Initialization Pattern in Hex field or use the default 00000000 value.
6. Click the **Confirm** button.

STOPPING, PAUSING OR RESUMING AN INITIALIZATION

To stop, pause or resume Initialization:

1. Click the **Administration** tab.

2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over Initialization and click the **Stop**, **Pause**, or **Resume** button.

REDUNDANCY CHECK ON A LOGICAL DRIVE

Redundancy Check is a routine maintenance procedure for fault-tolerant disk arrays (those with redundancy) that ensures all the data matches exactly. Redundancy Check can also correct inconsistencies.

To run **Redundancy Check** on a logical drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over **Redundancy Check** and click the **Start** button.
4. Check the boxes to the left of the logical drives you want to run.
5. Check the options you want:
 - **Auto Fix** – Attempts to repair the problem when it finds an error
 - **Pause on Error** – The process stops when it finds a non-repairable error
6. Click the **Confirm** button.

STOPPING, PAUSING OR RESUMING A REDUNDANCY CHECK

To stop, pause or resume **Redundancy Check**:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over Redundancy Check and click the **Stop**, **Pause**, or **Resume** button.

MIGRATING A LOGICAL DRIVE’S RAID LEVEL

The term “Migration” means either or both of the following:

- Change the RAID level of a logical drive.
- Expand the storage capacity of a logical drive.

Before you begin a migration, examine your current disk array to determine whether:

- The physical drives in your array can support the target RAID level.
- There is sufficient capacity to accommodate the target logical drive size.

If you need to add physical drives to your array, be sure there are unassigned physical drives installed in your RAID system before you begin migration.

See "Migration" on page 76, "RAID Levels" on page 184 and "RAID Level Migration" on page 189.

MIGRATING A LOGICAL DRIVE

To migrate a logical drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.
The list of background activities appears.
3. Mouse-over **Migrate** and click the **Start** button.
4. In the **Select Disk Array** drop-down menu, choose the source disk array.
5. In the **Select Physical Drives** diagram, click the drives to add them to your array.

The ID numbers of the chosen drives appear in the field below the diagram.

6. Click the **Next** button.
7. Check the box next to the logical drive you want to modify.
8. From the drop-down menu, choose a target RAID level.

The choice of RAID levels depends the number of physical drives in the disk array. See the Note below.

9. In the **Capacity** field, accept the current capacity.

Or check the **Expand Capacity** box and enter a greater capacity and size in MB, GB or TB.

If there is capacity remaining, you can create an additional logical drive.

10. Click the **Next** button.

The logical drive ID numbers, with the original and target RAID levels and capacities are shown

11. To accept the proposed target values, click the **Confirm** button.



Note

When you add physical drives to a RAID 10 array, it becomes a RAID 1E array by default. If you are adding an even number of physical drives to a RAID 10 array and you want the target array to be RAID 10, you must specify RAID 10 under RAID level.

CREATING A LUN CLONE

A LUN clone is an exact copy of the original LUN or logical drive, including all the data it contains, at one point in time. Use a LUN clone as a backup or to migrate a LUN from one system to another.



Important

The action of creating a LUN momentarily takes the original LUN or logical drive offline. While the logical drive is offline it is not available, i.e. no read or write operations are possible.

A LUN clone has the same capacity, stripe size, read and write policies as the original LUN. However, the LUN clone can be a different RAID level. The choice of RAID levels depends on the disk array. And if you have multiple disk arrays, you can create the LUN clone on a different disk array than the original LUN.

This action requires **Super User** or **Power User** privileges.

To create a LUN clone of a logical drive:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

The Logical Drive list appears.

3. Click the logical drive you want, then click the **LUN Clone** button.
4. Make settings as required:
 - From the **Choose a RAID level** drop-down menu, choose the RAID level of the LUN clone.
 - From the **Number of Copies** drop-down menu, choose the number of LUN clones you want to create. You can make up to 8 clones of a LUN at a time.
 - Check the box to the left of the Disk Array on which you want to create the LUN clone.
5. Click the **Next** button and review your choices.

6. Click the **Start** button to begin the cloning process.

The cloning progress bar displays.

Note the **Target Logical Drive ID**. Use this number to identify the LUN clone in the Logical Drive list.

If you chose a redundant RAID level, the LUN clone is automatically synchronized after creation.

After the LUN clone is created, you can manage it like any other logical drive. See Making Logical Drive Settings, Locating a Logical Drive, and Deleting a Logical Drive.

For users to access the LUN clone, you must map it to an initiator. See "Managing LUNs" on page 105.

MANAGING SPARE DRIVES

Spare drive management includes:

- "Viewing a List of Spare Drives" on page 101
- "Viewing Spare Drive Information" on page 102
- "Creating a Spare Drive Manually" on page 102
- "Deleting a Spare Drive" on page 102
- "Making Spare Drive Settings" on page 102
- "Locating a Spare Drive" on page 103
- "Running Spare Check" on page 103
- "Running a Transition on a Spare Drive" on page 103

VIEWING A LIST OF SPARE DRIVES

To view a list of spare drives:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

Spare Drive information displays, including:

- **ID** – Spare0, Spare1, etc.
- **Operational Status** – OK means normal
- **Configurable Capacity** – Usable capacity of the spare drive
- **Physical Drive ID** – ID number of the physical drive chosen for this spare
- **Revertible** – Yes or No
- **Spare Type** – Global or Dedicated
- **Dedicated to Array** – ID number of the disk array to which the spare is dedicated

VIEWING SPARE DRIVE INFORMATION

To view spare drive information:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

The list of spare drives appears.

3. Click the spare drive you want, then click the **View** button.

Spare Drive information displays, including:

- **Spare Drive ID** – Spare0, Spare1, etc.
- **Physical Drive ID** – ID number of the physical drive chosen for this spare
- **Operational Status** – OK means normal
- **Spare Type** – Global or Dedicated
- **Physical Capacity** – Total data capacity of the spare drive
- **Revertible** – Yes or No
- **Configurable Capacity** – Usable capacity of the spare drive
- **Spare Check Status** – Not Checked or Healthy
- **Media Patrol** – Enabled or Not Enabled
- **Dedicated to Array** – ID number of the disk array to which the spare is dedicated

For more information, see "Spare Drives" on page 193.

CREATING A SPARE DRIVE MANUALLY

This feature creates a spare drive only. You can also use the Wizard to create a disk array with logical drives and spare drives at the same time.

This action requires **Super User** or **Power User** privileges.

To create a spare drive:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.
3. Click the **Create Spare Drive** button.

4. For each of the following items, accept the default or change the settings as required:
 - Check the **Revertible** box if you want a revertible spare drive.
A revertible spare drive returns to its spare drive assignment after you replace the failed physical drive in the disk array and run the Transition function.
 - **Global** – Can be used by any disk array
 - **Dedicated to newly created disk array** – The disk array you are now creating.
5. In the **Select Physical Drives** diagram, click a drive to choose it for your spare.
The ID number for chosen drive appears in the field below the diagram.
6. Click the **Submit** button.
If you are done creating spare drives, click the **Finish** button.
To create another spare drive, click the **Create More** button.

DELETING A SPARE DRIVE

This action requires Administrator or a **Super User** privileges.

To delete a spare drive:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.
3. Click the spare drive you want, then click the **Delete** button.
4. In the Confirmation box, type the word “**confirm**” in the field provided and click the **Confirm** button.

MAKING SPARE DRIVE SETTINGS

To make spare drive settings:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.
3. Click the spare drive you want, then click the **Settings** button.
4. For each of the following items, accept the default or change the settings as required:
 - In the **Revertible** drop-down menu, choose Yes or No.
 - Check the **Media Patrol** box to enable Media Patrol on this spare drive. Uncheck to disable.
 - In the **Spare Type** drop-down menu, choose **Global** or **Dedicated**.
 - If you use chose a Dedicated spare, check the box beside the disk array to which this spare drive is assigned.
5. Click the **Save** button.

LOCATING A SPARE DRIVE

Spare drives are located in the same way as individual physical drives.

To locate a spare drive:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

The list of spare drives appears.

3. In the spare drive list, identify the physical drive ID number.
4. Click the **Device** tab.
5. Click the **Physical Drive** icon.

The list of physical drives appears.

6. Click the physical drive with the matching ID number and click the **Locate** button.

The drive carrier LED blinks for one minute.

RUNNING SPARE CHECK

Spare Check verifies the status of your spare drives.

To run spare check:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

The list of spare drives appears.

3. Click the spare drive you want, then click the **Spare Check** button.
4. Click the **Confirm** button.

After the “**Spare Check completed**” message appears, click the **View** button to see Spare Check Status.

RUNNING A TRANSITION ON A SPARE DRIVE

Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare. You must specify an unconfigured physical drive of the same or larger capacity and same media type as the revertible spare drive.

See "Transition" on page 77.

RUNNING A TRANSITION

To run a transition on a revertible spare drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over **Transition** and click the **Start** button.
4. From the **Source Physical Drive** drop-down menu, choose a Source disk array and the revertible spare drive.

Arrays have an ID No. The revertible spare has a Seq. No. (sequence number).

5. From the **Target Physical Drive** drop-down menu, choose a Target unconfigured drive.
6. Click the **Confirm** button.

STOPPING, PAUSING OR RESUMING A TRANSITION

To stop, pause or resume Transition:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of background activities appears.

3. Mouse-over **Transition** and click the **Stop**, **Pause**, or **Resume** button.

MANAGING INITIATORS

Initiator management includes:

- "Viewing a List of Initiators" on page 104
- "Adding an FC Initiator" on page 104
- "Deleting an FC Initiator" on page 104
- "Adding an iSCSI Initiator" on page 105

VIEWING A LIST OF INITIATORS

The Vess R2600’s initiator list displays initiators available for mapping to a LUN or logical drive. You must add initiators to the Vess R2600’s initiator list to make them available for mapping to a LUN.

To view a list of initiators:

1. Click the **Storage** tab.
2. Click the **Initiator** icon.

The list of initiators appears. Initiator information includes:

- **Index** – Initiator 0, Initiator 1, Initiator 2, etc.
- **Initiator Name**
 - Fibre Channel** – The World Wide Port Name of the initiator, composed of a series of eight, two-digit hexadecimal numbers.
 - iSCSI** – The iSCSI name of the initiator device, composed of a single text string.

ADDING AN FC INITIATOR

You must add an initiator to the Vess R2600’s initiator list in order to map your LUN or logical drive to the initiator.

METHOD 1: INPUTTING THE INITIATOR NAME

This action requires **Administrator** or **Super User** privileges.

To add a Fibre Channel initiator to the list:

1. Click the **Storage** tab.
2. Click the **Initiator** icon.

3. Click the **Add Initiator** button.
4. Input the initiator name in the fields provided.

An FC initiator name is the World Wide Port Name of the initiator, composed of a series of eight, two-digit hexadecimal numbers.

5. Click the **Submit** button.

The initiator is added.

METHOD 2: ADDING FROM A LIST

This action requires **Administrator** or **Super User** privileges.

To add a Fibre Channel initiator to the list:

1. Click the **Device** tab.
2. Click the **Fibre Channel Management** icon.
3. Click the Logged In **Device** tab.
4. Check the box next to the initiator you want to add.
5. Click the **Add to Initiator List** button.

The initiator is added, and its check box grays out.

DELETING AN FC INITIATOR



Caution

If you delete an initiator, you delete the LUN map associated with that initiator. Verify that the LUN map is no longer needed before deleting the initiator

This action requires **Administrator** or **Super User** privileges.

To delete an FC initiator:

1. Click the **Storage** tab.
2. Click the **Initiator** icon.
3. Click the initiator you want, then click the **Delete** button.
4. In the Confirmation box, type the word “**confirm**” in the field provided and click the **Confirm** button.

The initiator is removed from Vess R2600’s initiator list.

ADDING AN iSCSI INITIATOR

To add an iSCSI initiator to the list:

1. Click the **Storage** tab.
2. Click the **Initiator** icon.
3. Click the **Add Initiator** button.
4. Input the initiator name in the fields provided.

An iSCSI initiator name is the iSCSI name of the initiator device, composed of a single text string.

Example: **iqn.1991-05.com.microsoft:promise-29353b7**.

Obtain the initiator name from the initiator utility on your host system.

Note that the initiator name you input must match exactly in order for the connection to work.

5. Click the **Submit** button.

The initiator is added to the list.

MANAGING LUNs

LUN management includes:

- "Viewing a List of LUN Maps" on page 105
- "LUN Mapping and Masking" on page 105
- "Adding a LUN Map" on page 106
- "Editing a LUN Map" on page 106
- "Enabling and Disabling LUN Masking" on page 107

VIEWING A LIST OF LUN MAPS

To view a list of LUN maps:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.

The list of LUN maps appears.

LUN MAPPING AND MASKING

This feature applies to Fibre Channel and iSCSI subsystems and controls user access to storage resources.

- LUN Mapping – Maps LUNs to an initiator, so that the initiator can only access only the specified LUNs.
- LUN Masking – The process of applying a LUN Map.

To access LUN mapping:

1. Click the **Storage** tab.
2. Click the **LUN Masking & Mapping** icon.

On this screen, you can:

- Add an FC or iSCSI initiator to the Vess R2600’s initiator list.
- Enable LUN masking.
- Map a LUN to one or more initiators.

ADDING A LUN MAP

For FC systems or iSCSI systems, you can set up an Initiator LUN map.

A maximum of 256 logical drives can be mapped to an FC initiator or an iSCSI initiator.

To assign a LUN to an initiator, add the initiator first. See "Adding an FC Initiator" on page 104 or "Adding an iSCSI Initiator" on page 105.

LUN masking must be enabled in order to map a LUN. See “Enabling and Disabling LUN Masking” on page 107.

To add a LUN map:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.
3. Beside Active LUN Mapping Type,
 - FC subsystems, choose the **Initiator** option.
 - iSCSI subsystems, choose the **Initiator** option.

If you change the LUN Mapping Type, in the popup message type “**confirm**” and click the **Confirm** button.

4. Click the **LUN Mapping** button.

The first LUN Mapping screen appears.

This screen lets you choose initiators, ports, or targets, depending on the Active LUN Mapping Type.

5. Click the drop-down menu to choose the initiators, ports, or targets you want for the LUN map.

Choose your initiators, ports, or targets individually or choose all of them.

6. Click the **Next** button.

The second LUN Mapping screen appears.

7. Click a logical drive to highlight it. Then click the < button to assign the logical drive to an initiator or port.

Or click the << button to assign all logical drives to an initiator or port.

The logical drive moves to the Initiator, Port, or Target list with a default LUN of 0. Type the LUN you want to assign to this initiator, from 0 to 255.

Each logical drive can have only one unique LUN.

8. Click the **Next** button.

The final LUN Mapping screen appears showing the initiator or port and LUN map.

9. Click the **Submit** button.

The new LUN map is created.

EDITING A LUN MAP

Editing a LUN map is the action of assigning a logical drive or LUN to an initiator. By changing the assignment, you change the initiator’s access.

To edit a LUN map:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.

The list of LUN maps appears.

3. Click the LUN map you want to change, then click the **Setting** button.

4. Beside Active LUN Mapping Type,
 - FC subsystems, choose the **Initiator** option.
 - iSCSI subsystems, choose the **Initiator** option.

If you change the LUN Mapping Type, in the popup message type “confirm” and click the **Confirm** button.

5. Drag a logical drive from the Logical Drive list and drop it onto the Initiator list.
6. Click the **Next** button.

The LUN Mapping screen shows the edited LUN map.

7. Click the **Submit** button.

DELETING A LUN MAP

Deleting a LUN map prevents the initiator from accessing the LUN while LUN masking is enabled.

To delete a LUN map:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.

The list of LUN maps appears.

3. Click the LUN map you want, then click the **Delete** button.
4. In the Confirmation box, type the word “confirm” in the field provided and click the **Confirm** button.

ENABLING AND DISABLING LUN MASKING

LUN masking must be enabled in order to assign map your LUNs to your initiators and to use your existing LUN maps.

Disabling LUN masking allows all initiators to access all LUNs in your data storage. However, disabling LUN masking does not delete existing LUN maps.

These actions require **Administrator** or **Super User** privileges.

To enable or disable LUN masking:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.
3. Check the box to enable LUN Masking.

Or uncheck the box to disable LUN Masking.

LUN Masking starts or stops as soon as you make your setting.

MANAGING FIBRE CHANNEL CONNECTIONS

Fibre Channel management includes:

- "Viewing FC Node Information" on page 107
- "Viewing FC Port Information" on page 107
- "Making FC Port Settings" on page 108
- "Viewing FC Port Statistics" on page 108
- "Viewing a List of FC Initiators on the Fabric" on page 108
- "Viewing a List of FC SFPs" on page 109

Also see "Adding an FC Initiator" on page 104 and "Deleting an FC Initiator" on page 104.

VIEWING FC NODE INFORMATION

To view Fibre Channel node information:

1. Click the **Device** tab.
2. Click the **FC Management** icon.
3. Click the **Node** tab.

Node information includes:

- **Worldwide Node Name (WWNN)**
- **Maximum Frame Size**
- **Supported FC Class**
- **Supported speeds**

VIEWING FC PORT INFORMATION

To view Fibre Channel port information:

1. Click the **Device** tab.
2. Click the **FC Management** icon.
3. Click the **Port** tab.
4. Mouse-over an FC port to access and click the **View** button.

Port information includes:

<ul style="list-style-type: none">• FC Port ID - FC port number• Location - Controller ID number• Identifier - (hexadecimal)• State - Link status• Topology Attached -• Alias WWPN -• WWNN - World Wide Node Name	<ul style="list-style-type: none">• WWPN - Worldwide Port Name• Fabric WWNN• Fabric WWPN• Current Speed• Link Type -• Hard ALPA -
--	--

MAKING FC PORT SETTINGS

To make Fibre Channel port settings:

1. Click the **Device** tab.
2. Click the **FC Management** icon.
3. Click the **Port** tab.
4. Click the FC port you want to access and click the **Settings** button.
5. Make these changes as required:
 - Choose a configured link speed from the drop-down menu.
The choices are Auto (default), 2 Gb/s, 4 Gb/s, and 8 Gb/s.
 - Choose a topology from the drop-down menu.
 - **Enter** a Hard ALPA in the field provided.
Enter 255 to disable Hard ALPA.
6. Click the **Save** button.

PORT SETTING INFORMATION

The table below shows the type of attached topology you achieve based on your connection type and the configured topology you select.

Example 1: If you connect the Vess R2600 to an FC switch and choose NL-Port topology, you create a Public Loop attached topology.

Example 2: If you have a Point-to-Point attached topology, you made a direct connection (no FC switch) and selected N-port topology.



Note

In some cases, HBA settings to N-Port only work if connected to the switch. Refer to your HBA manual for more information.

VIEWING FC PORT STATISTICS

To view Fibre Channel port statistics:

1. Click the **Device** tab.
2. Click the **FC Management** icon.
3. Click the **Statistics** tab.
4. Mouse over the FC port you want to access and click the **View** button.

To clear FC port statistics, see "Clearing Statistics" on page 55.

VIEWING A LIST OF FC INITIATORS ON THE FABRIC

To view a list Fibre Channel initiators on the fabric:

1. Click the **Device** tab.
2. Click the **FC Management** icon.
3. Click the **Initiators on Fabric** tab.

Also see "Viewing a List of Initiators" on page 104.

VIEWING A LIST OF FC LOGGED-IN DEVICES

Logged-in devices refers to all Fibre Channel devices currently logged into the Vess R2600. The device list includes:

- **FC ports**
- **FC switches, if attached**
- **FC initiators**

To view a list FC logged-in devices:

1. Click the **Device** tab.
2. Click the **FC Management** icon.
3. Click the **Logged In Device** tab.

VIEWING A LIST OF FC SFPs

The term SFP refers to Small Form Pluggable transceivers used in Fibre Channel ports. The SFPs convert electrical signals to optical signals and send them over the Fibre Channel fabric, where another transceiver converts the optical signal back to an electrical signal again.

To view a list FC SFPs:

1. Click the **Device** tab.
2. Click the **FC Management** icon.
3. Click the **SFP** tab.

SFP information includes:

- **FC port ID**
- **Controller ID**
- **Connector type**
- **Transceiver type**
- **Transceiver code**
- **Vendor name**

MANAGING iSCSI CONNECTIONS

iSCSI management includes:

- "Viewing a List of iSCSI Targets" on page 110
- "Viewing iSCSI Target Information" on page 110
- "Making iSCSI Target Settings" on page 111
- "Viewing a List of iSCSI Portals" on page 111
- "Viewing a List of iSCSI Portals" on page 111
- "Adding iSCSI Portals" on page 111
- "Making iSCSI Portal Settings" on page 112
- "Deleting iSCSI Portals" on page 112
- "Viewing a List of iSCSI Ports" on page 112
- "Making iSCSI Port Settings" on page 113
- "Viewing a List of iSCSI Trunks" on page 113
- "Making iSCSI Trunk Settings" on page 114
- "Deleting iSCSI Trunks" on page 114
- "Viewing iSCSI Session Information" on page 114
- "Deleting an iSCSI Session" on page 115
- "Making iSCSI iSNS Settings" on page 115
- "Viewing a List of iSCSI CHAPs" on page 116
- "Adding iSCSI CHAPs" on page 116
- "Deleting iSCSI CHAPs" on page 116

VIEWING A LIST OF iSCSI TARGETS

A **target** is a logical device on the Vess R2600 subsystem. The default target exposes all logical drives and is associated with all portals on the subsystem.

To view a list of iSCSI targets:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Target** tab.

Target information includes:

- **ID** – ID number of the target
- **Alias** – If assigned
- **Assigned Portals** – Portals assigned under this target

VIEWING iSCSI TARGET INFORMATION

To view information about a target:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Target** tab.
4. Click the target you want, then click the **View** button.

Target information includes:

- **ID** – ID number of the target.
- **Name** – iSCSI qualified name (iqn) of this target.
- **Alias** – Maximum of 32 characters. Use letters, numbers, space between words, and underscore. An alias is optional.*
- **Status** – Up or down.
- **Error Recovery Level** – Error recovery level supported.
- **Initial R2T** – Allows initiator to begin sending data to a target without receiving a ready to transfer command.
- **Max Outstanding R2T** – Maximum number of R2T PDUs the target can have outstanding for a single iSCSI command.
- **Max Burst Length** – Maximum length of a solicited data sequence in bytes.
- **Data Digest** – Adds a data digest (CRC).*
- **Header Digest** – Enables the use of header digest (CRC).*
- **Data Sequence in Order** – Enables placement of data in sequence order
- **Data PTU in Order** – Enables placement of data in PDU order
- **Default Time to Wait** – After a dropped connection, the number of seconds to wait before attempting to reconnect
- **Default Time to Retain** – Number of seconds after time to wait (above) before reassigning outstanding commands
- **Uni-directional CHAP Authentication** – Uni-directional (peer) CHAP authentication, enabled or disabled*
- **Bi-directional CHAP Authentication** – Bi-directional (local) CHAP authentication, enabled or disabled*
- **Maximum Connections** – The maximum number of concurrent connections
- **Immediate Data** – Enables the initiator to send unsolicited data with the iSCSI command PDU.
- **First Burst Length** – In bytes.
- **NOP-In** - Check iSCSI connection status*

Items marked with an asterisk (*) are adjustable under "Making iSCSI Target Settings"

MAKING iSCSI TARGET SETTINGS

To make target settings:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Target** tab.
4. Click the target you want, then click the **Settings** button.
5. Make settings changes are required:
 - **Alias**
 - **Enable Header Digest**
 - **Enable Data Digest**
 - **Enable Uni-directional CHAP Authentication**
 - **Enable Bi-directional CHAP Authentication**
 - **Enable NOP-In**
6. Click the **Submit** button.

VIEWING A LIST OF iSCSI PORTALS

To view a list of iSCSI portals:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Portal** tab.

Portal information includes:

- **ID** – Portal number. Starts at 0.
- **IP Address** – IP address of the portal.
- **Controller ID** – RAID controller ID, 1 or 2.
- **Port ID** – Physical port on the RAID controller, 1 to 4.
- **Trunk ID** – Trunk ID, 1 to 8. Refers to portals associated with a trunk (link aggregation). N/A means this portal is not associated with a trunk.
- **VLAN Tag** – VLAN Tag, 1 to 4094. Refers to portals associated with a Virtual Local Area Network (VLAN). N/A means this portal is not associated with a VLAN.

VIEWING iSCSI PORTAL INFORMATION

To view information about a portal:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Portal** tab.
4. Click the portal you want, then click the **View** button.

Portal information includes:

- **Portal ID** – Portal number. Starts at 0.
- **Trunk ID** – 1 to 8. Refers to portals associated with a trunk (link aggregation). N/A means this portal is not associated with a trunk.
- **Controller ID** – RAID controller ID, 1 or 2.
- **VLAN Tag** – 1 to 4094. Refers to portals associated with a Virtual Local Area Network (VLAN). N/A means this portal is not associated with a VLAN.
- **PrimaryIP** - IP Address
- **PrimaryIPMask** - subnet mask
- **Gateway** - IP address of the gateway routing device
- **Port ID** – Physical port on the RAID controller, 1 to 4.
- **Interface Name** – Ethernet interface names.
- **Associated Type** – PHY, VLAN, or Trunk.
- **DHCP** – Enabled or disabled.* DHCP is currently supported only for IPv4.
- **TCP Port Number** – TCP port number. 3260 is the default and recommended number.
- **Assigned Targets** – IDs of the targets to which this Portal is assigned. N/A means no target is assigned.
See "Adding iSCSI Portals" on page 111

Items marked with an asterisk (*) are adjustable under "Making iSCSI Portal Settings"

ADDING iSCSI PORTALS

Vess R2600 supports up to 32 iSCSI portals. Each iSCSI portal can belong to a different VLAN for a maximum of 32 VLANs.

If you plan to associate the new portal with a trunk, create the trunk first. See “Adding iSCSI Trunks” on page 113.

For more information about iSCSI VLANs, see "iSCSI on a VLAN" on page 197.

To add a portal:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Portal** tab.
4. Click the **Create Portal** button.
5. Make your choices and inputs as required:
 - Choose an Association type from the option list.
The choices are ***PHY*** or ***Trunk***,
 - If you are creating a ***PHY*** association, choose
 - **Controller ID** (1 or 2) from the drop-down menu.
 - Choose a **Port ID** (1 to 4) from the drop-down menu.
 - If you want to use VLAN, click **enable VLAN** and fill a VLAN tag from (1 to 4094).
 - If you are creating a ***Trunk*** association, choose a **Trunk ID** (1 to 8) from the drop-down menu.
 - Type the IP address of the portal in the field provided.
 - Type the subnet mask of the portal in the field provided.
 - Type the gateway IP address of the portal in the field provided.
 - From the IP Type drop-down menu, choose IPv4 or IPv6.
DHCP is currently supported only for IPv4.

6. Click the **Submit** button.

The new portal is added to the list.

To assign a portal to a target, see "Assigning a Portal to an iSCSI Target"

MAKING iSCSI PORTAL SETTINGS

To make iSCSI portal settings:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Portal** tab.
4. Click the portal you want, then click the **Settings** button.
5. Make settings changes as needed:
 - If you have a ***Trunk*** association, choose a Trunk ID (1 to 8) from the drop-down menu.
 - Type the IP address of the portal in the field provided.
 - Type the subnet mask of the portal in the field provided.
 - If you have a ***VLAN*** association, enter a VLAN tag (1 to 4094) in the field provided.

- From the IP Type drop-down menu, choose IPv4 or IPv6.
DHCP is currently supported only for IPv4.

6. Click the **Submit** button.

DELETING iSCSI PORTALS

To delete an iSCSI portal:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Portal** tab.
4. Click the portal you want, then click the **Delete** button.
5. In the **Confirmation** box, type the word “**confirm**” in the field provided and click the **Confirm** button.

The portal is removed from the list.

VIEWING A LIST OF iSCSI PORTS

An iSCSI port is the physical iSCSI connection on the Vess R2600. There are four iSCSI ports on each RAID controller for a total of eight per subsystem.

To view a list of ports:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Port** tab.

Port information includes:

- **Port ID** – ID number of the port
- **Controller ID** – 1 or 2
- **Link Status** – Up or down, active or Inactive
- **Jumbo Frames** – Enabled or disabled*
- **Current Speed** – In Mb/s
- **Assigned Portals** – Portals to which this port is assigned

Items marked with an asterisk (*) are adjustable under “Making iSCSI Port Settings”

VIEWING iSCSI PORT INFORMATION

To view information about a port:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Port** tab.
4. Click the port you want, then click the **View** button.

Port information includes:

- **Controller ID** – ID of the RAID controller where the port is located
- **Status** – Enabled or disabled*
- **Jumbo Frames** – Enabled or disabled*
- **Link Status** – Up or down, active or inactive
- **MAC Address** – MAC address of the target port
- **Max Supported Speed** – Maximum speed supported (1 Gb/s)
- **Current Speed** – Current or actual speed of the target port
- **Relative Portals** – The portals corresponding to this target port

Items marked with an asterisk (*) are adjustable under "Making iSCSI Port Settings"

MAKING iSCSI PORT SETTINGS

To make iSCSI port settings:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Port** tab.
4. Click the port you want, then click the **Settings** button.
5. Make settings changes as required:
 - **Jumbo Frames** – Check to enable jumbo frame support on this port. Uncheck to disable.
6. Click the **Submit** button.

VIEWING A LIST OF iSCSI TRUNKS

A trunk is the aggregation of two or more iSCSI ports to increase bandwidth.

To view a list of trunks:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Trunk** tab.

Trunk information includes:

- **Trunk ID** – ID number of the trunk
- **Controller ID** – ID of the RAID controller, 1 or 2
- **Master Port** – ID of the master port
- **Slave Ports** – IDs of the slave ports
- **Failed Ports** – IDs of any ports that are not working
- **State** – Optimal, Sub-Optimal, or Failed

Failed ports result in sub-optimal and failed trunks.

ADDING iSCSI TRUNKS

Ports must be ***enabled*** to add them to a trunk. Vess R2600 supports a maximum of eight trunks.

You cannot use an iSCSI port that has portals configured to it. See "Viewing a List of iSCSI Portals" and "Deleting iSCSI Portals."

To add an iSCSI trunk:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Trunk** tab.
4. Click the **Create Trunk** button.
5. Make your choices as required:
 - **Controller ID** – ID of the RAID controller, 1 or 2
 - **Master Port number** – ID of the master port
 - **Slave Port number** – IDs the slave ports

6. Click the Submit button.

The new trunk is added to the list.

Specify the trunk when your create a portal. See "Adding iSCSI Portals"

MAKING iSCSI TRUNK SETTINGS

To make trunk settings:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Trunk** tab.
4. Click the trunk you want, then click the **Settings** button.
5. Make changes as required:
 - **Controller ID** – ID of the RAID controller, 1 or 2
 - **Master Port number** – ID of the master port
 - **Slave Port number** – IDs the slave ports
6. Click the **Submit** button.

DELETING iSCSI TRUNKS

Before you can delete a trunk, you must delete any portals configured on it. See "Deleting iSCSI Portals" on page 112.

To delete an iSCSI trunk:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Trunk** tab.
4. Click the trunk you want, then click the **Delete** button.
5. In the **Confirmation** box, type the word “**confirm**” in the field provided and click the **Confirm** button.

The trunk is removed from the list.

VIEWING A LIST OF iSCSI SESSIONS

To view a list of iSCSI sessions:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Session** tab.

iSCSI session information includes:

- **ID** – ID number of the session
- **Target Name** – Alias of the target
- **Initiator Name** – Part of the IQN
- **Portal ID** – ID number of the portal
- **Status** – Active or inactive.

VIEWING iSCSI SESSION INFORMATION

To view a list of iSCSI sessions:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Session** tab.
4. Click the iSCSI session you want and click the View button.

Information includes:

- **Session ID** – ID number of the session
- **Status** – Active or inactive
- **Target Alias**
- **Initiator Name** – iSCSI qualified name (iqn)
- **Portal IP** – IP address of the portal
- **Device Type** – Initiator or target
- **Target Portal Group** – ID number
- **TSIH** – Target session identifying handle
- **Execution Throttle** – Max number of outstanding commands on any one port
- **Max Outstanding R2T** – Number of PDUs ready to transfer
- **Default Time to Retain** – In seconds
- **Max Burst Length** – In bytes
- **Initial R2T** – Enabled or disabled
- **Data Digest** – Enabled or disabled
- **Data PDU in Order** – Enabled or disabled
- **Portal ID** – ID number of the portal
- **Keep Alive** – Enabled or disabled
- **Target Name** – iSCSI qualified name (iqn)
- **Initiator IP** – IP address of the initiator
- **Device Access Control** – Enabled or disabled
- **Initiator Source Port** – ID number
- **ISID** – Initiator session ID number
- **Max Rcv Data Seg Length** – Receive data segment length
- **First Burst Length** – In bytes
- **Default Time to Wait** – In seconds
- **Immediate Data** – Enabled or disabled
- **Header Digest** – Enabled or disabled
- **CHAP Authentication Type** – None, Local, Peer
- **Data Seq in Order** – Enabled or disabled

DELETING AN iSCSI SESSION

To delete an iSCSI session:

1. Click the Device tab.
2. Click the **iSCSI Management** icon.
3. Click the **Session** tab.
4. Click the iSCSI session you want and click the **Delete** button.
5. Type “**confirm**” in the field provided, then click the **Confirm** button.

VIEWING iSCSI iSNS INFORMATION

To view information about iSNS:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the iSNS tab.

The information includes:

- **Auto iSNS IP** – Yes means the IP address is assigned automatically
- **iSNS Enabled** – Yes means the iSNS feature is enabled*
- **iSNS Server IP Address** – IP address of the iSNS Server*
- **iSNS Port** – 3205 is the default and recommended value*

Items marked with an asterisk (*) are adjustable under Making iSCSI iSNS Settings.

MAKING iSCSI iSNS SETTINGS

To make iSNS settings:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **iSNS** tab.
4. Click the **iSNS Settings** button.
5. Make settings changes are required:
 - Check the box to enable iSNS. Uncheck to disable.
 - **Enter** the iSNS server IP address.
 - **Enter** a new iSNS Port number. The range is 1 to 65535.
6. Click the **Submit** button.

VIEWING A LIST OF iSCSI CHAPs

To view a list of iSCSI CHAPs:

1. Click the Device tab.
2. Click the **iSCSI Management** icon.
3. Click the CHAP tab.

CHAP information includes:

- **Index** – ID number of the CHAP
- **Name** – User assigned name of the CHAP
- **Type** – Peer or local
Peer is one-way or uni-directional.
Local is two-way or bi-directional.
- **Target ID** – ID number of the target (logical drive) where the CHAP is used. N/A means that no target is assigned.

ADDING iSCSI CHAPs

To add an iSCSI CHAP:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **CHAP** tab.
4. Click the **Create CHAP** button.
5. Make your choices and inputs as required:
 - **Enter** a name in the Name field.
 - Choose a CHAP type.
Peer is one-way or uni-directional.
Local is two-way or bi-directional.
 - **Enter** a secret of 16 characters in the Secret field.
 - **Enter** the secret again in the Retype Secret field.

6. Click the **Submit** button.

The new CHAP is added to the list.

MAKING iSCSI CHAP SETTINGS

When you change CHAP settings, you must change the secret. You cannot change the type (peer or local).

To make iSCSI CHAP settings:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **CHAP** tab.
4. Click the CHAP you want, then click the **Settings** button.
5. Make settings changes are required:
 - **Enter** a name in the **Name** field.
 - **Enter** the current secret in the **Current Secret** field.
 - **Enter** a new secret of 12 or more characters in the **Secret** field.
 - **Enter** the new secret again in the **Retype Secret** field.

6. Click the **Submit** button.

DELETING iSCSI CHAPs

To delete an iSCSI CHAP:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **CHAP** tab.
4. Click the CHAP you want, then click the **Delete** button.
5. Click the **Confirm** button.

The CHAP is removed from the list.

PINGING A HOST OR SERVER ON THE iSCSI NETWORK

This function enables you to ping other network nodes through any one of the Vess R2600’s iSCSI ports.

To ping a host or server on the network:

1. Click the **Device** tab.
2. Click the **iSCSI Management** icon.
3. Click the **Ping** tab.
4. Type the IP address of the host or server into the IP Address field.
5. Choose the port Type from the drop-down menu.
 - iSCSI means an iSCSI port
 - Mgmt means the Vess R2600’s virtual management port
6. If you chose iSCSI port, choose the RAID controller and port number from the drop-down menus.
7. Type the number of packets you want to send in the **Number of Packets to Ping** field.

Four packets are commonly used for a ping.

8. Click the **Start** button.

In a few moments, the result displays under the **Device** tab as

Ping succeeded or ***Ping failed***.

CHAPTER 5: MANAGEMENT WITH THE CLU

This chapter covers the following topics:

- "Managing the Subsystem (CLU)" on page 122
- "Managing the RAID Controllers (CLU)" on page 124
- "Managing the Enclosure (CLU)" on page 126
- "Managing Disk Arrays (CLU)" on page 132
- "Managing Spare Drives (CLU)" on page 137
- "Managing Logical Drives (CLU)" on page 138
- "Managing the Network Connection (CLU)" on page 142
- "Managing Fibre Channel Connections (CLU)" on page 143
- "Managing iSCSI Connections (CLU)" on page 146
- "Managing Background Activity (CLU)" on page 154
- "Working with the Event Viewer (CLU)" on page 155
- "Working with LUN Mapping (CLU)" on page 156
- "Managing UPS Units (CLU)" on page 159
- "Managing Users (CLU)" on page 160
- "Working with Software Management (CLU)" on page 163
- "Flashing through TFTP" on page 166
- "Viewing Flash Image Information (CLU)" on page 166
- "Clearing Statistics (CLU)" on page 167
- "Restoring Factory Defaults (CLU)" on page 167
- Shutting Down the Subsystem (CLU)" on page 168
- "Starting Up After Shutdown" on page 169
- "Restarting the Subsystem" on page 170
- "Buzzer" on page 171

For information about the Vess R2600 audible alarm and LEDs, see "Chapter 8: Troubleshooting" on page 201.

INITIAL CONNECTION


Making an initial connection includes the following functions:

- "Making a Serial Connection" on page 119
- "Making a SSH Connection" on page 120
- "Logging Into the CLI" on page 120
- "Accessing Online Help" on page 121
- "Exiting the CLU" on page 121
- "Logging Out of the CLI" on page 121
- "Logging Back Into the CLI and CLU" on page 121

MAKING A SERIAL CONNECTION

Before you begin, be sure the RJ-11-to-DB9 serial data cable is connected between the Host PC and Vess R2600, and that both machines are booted and running.

Figure 5 - 1: Serial ports on the controllers (Vess R2600fiD)




Then do the following actions:

1. Change your terminal emulation program settings to match the following specifications:
 - Bits per second: 115200
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: none
2. Start your PC's terminal VT100 or ANSI emulation program.
3. Press **Enter** once to launch the CLI.

MAKING A TELNET CONNECTION

A Telnet connection requires a network connection between the Host PC and the Management (Ethernet) port on the Vess R2600 controller.

Figure 5 - 2: Management port on the RAID controller



To start the telnet program:

1. Go to the command line prompt (Windows) or click the terminal icon (Linux).
2. Type **telnet 10.0.0.1** and press **Enter**.

The IP address above is only an example.

Use the Management port IP address of your Vess R2600.

The Telnet default port number is 2300.

3. Press **Enter** once to launch the CLI.

MAKING A SSH CONNECTION

A Secure Shell (SSH) connection requires a network connection between the Host PC and the Management (Ethernet) port on the Vess R2600 controller.

See above,Management port on the RAID controller.

Windows PCs require you to install a SSH application on the PC.

WINDOWS

To start the Windows SSH program:

1. Open the SSH application from the Start menu.
2. **Enter** the IP address and SSH port number of the Vess R2600 in the fields provided.

The SSH default port number is 22.

3. Press **Enter** once to launch the CLI.

LINUX

To start the Linux SSH program:

1. Click the terminal icon.
2. Type ssh administrator@10.0.0.1 and press Enter.

The IP address above is only an example.

Use the Management port IP address of your Vess R2600.

The SSH default port number is 22.

3. Press **Enter** once to launch the CLI.

LOGGING INTO THE CLI

1. At the Login prompt, type the user name and press **Enter**.

The default user name is **administrator**.

2. At the Password prompt, type the password and press **Enter**.

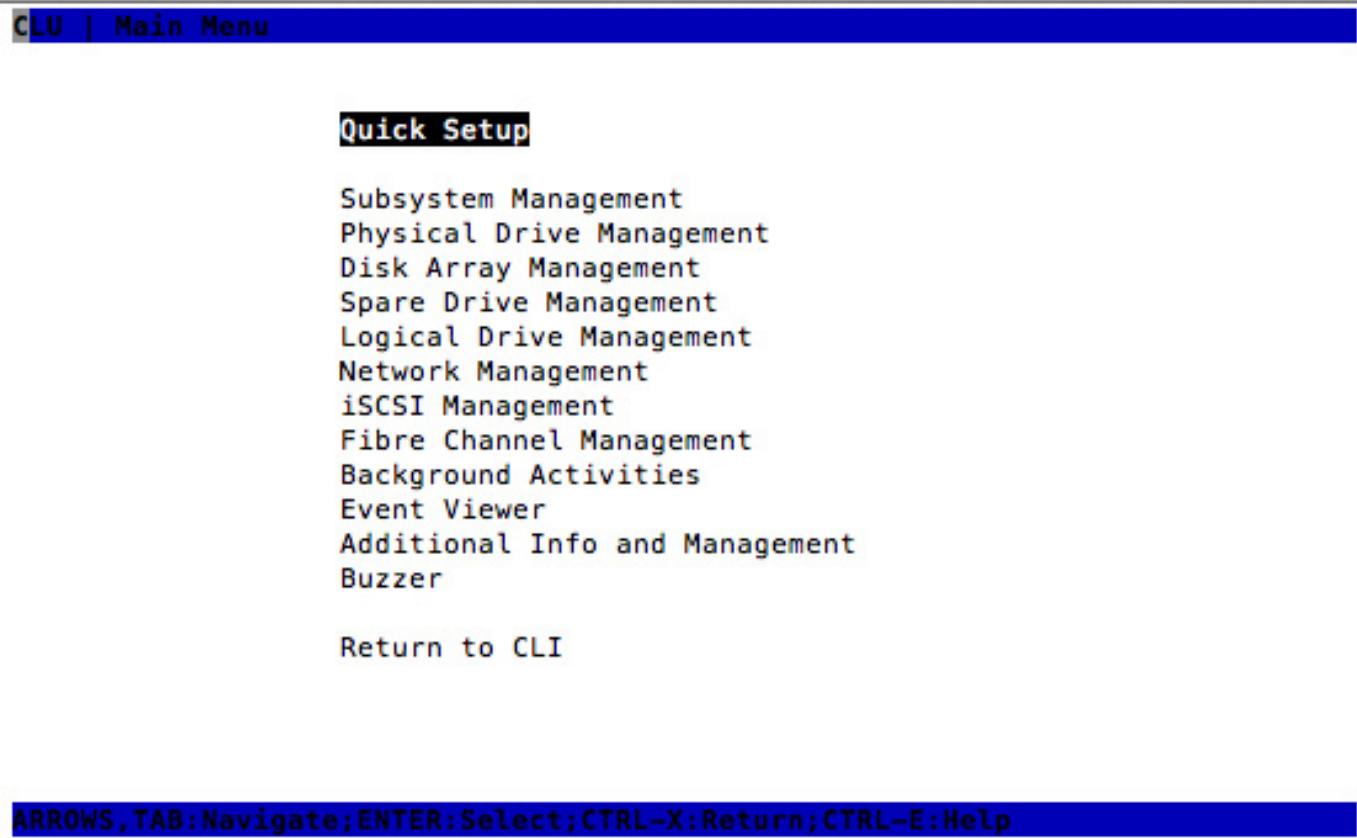
The default password is **password**.

The CLI screen appears.

3. At the **administrator@cli>** prompt, type **menu** and press **Enter**.

The CLU Main Menu appears.

Figure 5 - 3: CLU main menu



Quick Setup— A sequence of four steps to setup system date and time, Management port, and RAID configuration. See "Setting-up Vess R2600 with the CLU" on page 41.

Subsystem Management – Subsystem settings, Controller settings, statistics, lock/unlock the subsystem, set date and time, Enclosure settings, FRUs and Topology.

Physical Drive Management – Assign an alias, force a physical drive offline or online, clear a Stale or PFD condition, change global physical drive settings, and locate a physical drive.

Disk Array Management – Assign an alias, view array information, create and delete disk arrays, transport, rebuild, PDM, and transition functions, accept and incomplete array, locate a disk array, create, and delete logical drives.

Spare Drive Management – View a list of spare drives, create, modify, and delete spare drives, and run Spare Check.

Logical Drive Management – Assign an alias, set cache policies, view logical drive information, run initialization and Redundancy Check, create a LUN clone, and locate a logical drive.

Network Management – Set IP addresses for Virtual and Maintenance Mode Ports, gateway, and DNS server; subnet mask.

Fibre Channel Management – Node information, Port information, settings, SFPs, and statistics, Logged-in devices, add initiator to the list.

iSCSI Management – Targets, Ports, Portals, Sessions, iSNS options, CHAPs, Ping, Trunks, Logged-in devices, add initiator to the list.

Background Activities – Summary of running and scheduled activity, settings for Media Patrol, Auto Rebuild, Rebuild, Migration, PDM, Transition, Synchronization, Initialization, Redundancy Check rate, and thresholds.

Event Viewer – View runtime and NVRAM event logs.

Additional Info and Management – LUN mapping, UPS management, User management, Software services management, Flash through TFTP (Firmware update), Clear Statistics, Restore Default Settings, Shutdown or Restart the subsystem.

Buzzer – Enable, disable or silence the buzzer (audible alarm).

ACCESSING ONLINE HELP

To access online help on any CLU screen, press **Control-AE**.

To return to the CLU, press **Enter**.

EXITING THE CLU

1. Highlight **Return to Previous Menu** and press **Enter**.

Repeat this action until you arrive at the Main Menu.
2. From the Main Menu, highlight **Return to CLI** and press **Enter** to exit
3. Close the terminal emulation, Telnet, SSH, or terminal window.

LOGGING OUT OF THE CLI

When you shut down or restart the Vess R2600 subsystem, you are automatically logged out of the CLI.

To manually log out of the CLI (no shut down or restart):

At the **username@cli>** prompt, type logout and press **Enter**.

The prompt changes to **cli>**.

LOGGING BACK INTO THE CLI AND CLU

To log into the CLI and CLU after a manual logout:

1. At the **cli:>** prompt, type login followed by your user name and press **Enter**.
2. At the Password: prompt, type your password and press **Enter**.
3. At the **username@cli>** prompt, type **menu** and press **Enter** to open the CLU.

MANAGING THE SUBSYSTEM (CLU)

Subsystem Management includes the following functions:

- "Making Subsystem Settings (CLU)" on page 122
- "Locking or Unlocking the Subsystem (CLU)" on page 122
- "Setting Subsystem Date and Time (CLU)" on page 123
- "Making NTP Settings (CLU)" on page 123
- "Synchronizing with a NTP Server (CLU)" on page 123

MAKING SUBSYSTEM SETTINGS (CLU)

An alias is optional. To set an Alias for this subsystem:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Subsystem Settings** and press **Enter**.
3. Make changes as required:
 - Type and alias into the Alias field.
Maximum of 48 characters. Use letters, numbers, space between words and underscore.
 - Highlight **Redundancy Type** and press the spacebar to toggle between Active-Active and Active-Standby.
Active-Active – Both RAID controllers are active and can share the load
Active-Standby – One RAID controller is in standby mode and goes active if the other fails
 - Highlight **Cache Mirroring** and press the spacebar to toggle between Enabled and Disabled.
4. Press **Control-A** to save your settings.

RUNNING MEDIA PATROL (CLU)

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. Media Patrol checks all physical drives assigned to disk arrays and spare drives. It does not check unconfigured drives.

To start, stop, pause or resume Media Patrol:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Media Patrol** and press **Enter**.
3. Highlight **Start, Stop, Pause, or Resume** and press **Enter**.
4. If you chose **Stop**, press **Y** to confirm.

LOCKING OR UNLOCKING THE SUBSYSTEM (CLU)

The lock prevents other sessions (including sessions with the same user) from making a configuration change to the controller until the lock expires or a forced unlock is done. When the user who locked the controller logs out, the lock is automatically released.

SETTING THE LOCK

To set the lock:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Lock Management** and press **Enter**.
3. In the Lock Time field, type a lock time in minutes.

1440 minutes = 24 hours

4. Highlight **Lock** and press **Enter**.

RESETTING THE LOCK

To reset the lock with a new time:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Lock Management** and press **Enter**.
3. In the Lock Time field, type a lock time in minutes.

1 to 1440 minutes (24 hours)

4. Highlight **Renew** and press **Enter**.

RELEASING THE LOCK

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Lock Management** and press **Enter**.
3. Highlight **Unlock** and press **Enter**.

RELEASING A LOCK SET BY ANOTHER USER

To release somebody else’s lock:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Lock Management** and press **Enter**.
3. Highlight **Force Unlock** and press the Spacebar to change to **Yes**.
4. Highlight **Unlock** and press **Enter**.

SETTING SUBSYSTEM DATE AND TIME (CLU)

Use this screen to make Date and Time settings:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Modify System Date & Time** and press **Enter**.
3. Highlight the **System Date** or **System Time** setting.
4. Press the backspace key to erase the current value.
5. Type in a new value.
6. Press **Control-A** to save your settings.

MAKING NTP SETTINGS (CLU)

After you have made Network Time Protocol (NTP) settings, the Vess R2600 subsystem synchronizes with a NTP server.

- At startup
- Every night
- When you synchronize manually

To make NTP settings for the subsystem:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **NTP Management** and press **Enter**.

3. Highlight **NTP Settings** and press **Enter**.
4. Make the following settings as required:
 - Highlight **NTP Service** and press the spacebar to toggle between **Enabled** and **Disabled**.
 - Highlight **Time Server (1)**, **Time Server (2)**, or **Time Server (3)** and type a server name.
Example: 0.us.pool.ntp.org
You can have up to 3 NTP servers.
 - Highlight **Time Zone** and press the spacebar to toggle through GMT, GMT+, and GMT-.
For GMT+ and GMT-, type the hour from 0:00 to 13:00 GMT for your time zone.
 - Highlight **Daylight Savings Time** and press the spacebar to toggle between **Enable** and **Disable**.
If Daylight Savings Time is Enabled, highlight the **Start Month** and **End Month** and enter a number from 1 to 12.

Then highlight the **Week** and **Day** and toggle to make your choices.
5. Press Control-A to save your settings.



Notes

The NTP server name shown is an example only. You must find and enter your local NTP server name. GMT is the older designation for UTC.

SYNCHRONIZING WITH A NTP SERVER (CLU)

The Vess R2600 subsystem automatically synchronizes with a NTP server every night and a startup. You have the option of synchronizing manually at any time.

To manually synchronize the Vess R2600 with a NTP server:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **NTP Management** and press **Enter**.
3. Highlight **Start Time Sync** and press **Enter**.
4. Press Y to confirm.

To verify, check Last Synchronization Time and Last Synchronization Result.

MANAGING THE RAID CONTROLLERS (CLU)

RAID controller management includes the following functions:

- "Viewing Controller Information (CLU)" on page 124
- "Making Controller Settings (CLU)" on page 124
- "Locating the Controller (CLU)" on page 125

VIEWING CONTROLLER INFORMATION (CLU)

Controller Management includes information, settings and statistics.

To access Controller Management:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Controller Management** and press **Enter**.

The Controller summary information includes:

- **Controller ID** – 1 or 2
- **Alias** – if assigned
- **Operational Status** – OK means normal. Might show BGA running. Not present indicates a malfunction or no controller is installed
- **Readiness Status** – Active or Standby is normal. N/A means not accessible

3. Highlight the controller you want and press **Enter**.

To access additional controller information, highlight **Advanced Information** and press **Enter**.

To access controller statistics, highlight **Controller Statistics** and press **Enter**.

CLEARING STATISTICS

To clear controller statistics, see "Clearing Statistics (CLU)" on page 167.

CLEARING AN ORPHAN WATERMARK (CLU)

This condition is the result of a disk drive failure during an NVRAM RAID level migration on a disk array.

To clear an orphan watermark:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Controller Management** and press **Enter**.
3. Highlight one of the controllers and press **Enter**.
4. Highlight **Clear Orphan Watermark** and press **Enter**.

The condition is cleared. See "Physical Drive Problems" on page 213 for more information.

MAKING CONTROLLER SETTINGS (CLU)

If your subsystem has two controllers, any settings you make to one controller automatically apply to the other controller.

To make Controller settings:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Controller Management** and press **Enter**.
3. Highlight the controller you want and press **Enter**.
4. Highlight **Controller Settings** and press **Enter**.
5. Make the following settings as required:
 - Type an alias into the Alias field.
Maximum of 48 characters. Use letters, numbers, space between words and underscore. An alias is optional.
 - Highlight **LUN Affinity** and press the spacebar to toggle between **Enabled** and **Disabled**.
RAID controllers must be set to **Active-Active**. See "Making Subsystem Settings (CLU)" on page 122 and "LUN Affinity" on page 194.
 - Highlight **Coercion** and press the spacebar to toggle between **Enabled** and **Disabled**.
For more information, see "Capacity Coercion" on page 196.
 - Highlight **Coercion Method** and press the spacebar to toggle through:
 - GB Truncate** – Reduces the capacity to the nearest 1 GB boundary.
 - 10 GB Truncate** – Reduces the capacity to the nearest 10 GB boundary.
 - Grp (group) Rounding** – Uses an algorithm to determine truncation. Results in the maximum amount of usable drive capacity.
 - Table Rounding** – Applies a predefined table to determine truncation.

- Highlight **Host Cache Flushing** and press the spacebar to toggle between **Enable** and **Disable**.
For more information, see "Host Cache Flushing" on page 195.
- Highlight **Cache Flush Interval** and press the backspace key to erase the current value. Type a new interval value.
The range is 1 to 12 seconds. For more information, see "Cache Policy" on page 195.
- Highlight **SMART** and press the spacebar to toggle between **Enable** and **Disable**.
- Highlight **SMART Poll Interval** and press the backspace key to erase the current value. Type a new interval value (1 to 1440 minutes).
- Highlight **Poll Interval** and press the backspace key to erase the current value. Type a new interval value (15 to 255 seconds).
- Highlight **Adaptive Writeback Cache** and press the spacebar to toggle between **Enabled** and **Disabled**.
For more information, see "Host Cache Flushing" on page 195.
- Highlight **Forced Read Ahead Cache** and press the spacebar to toggle between **Enabled** and **Disabled**.
For more information, see "Forced Read-Ahead Cache" on page 195.
- Highlight **HDD Power Saving** and the spacebar to choose a time period.
After an HDD has been idle for a set period of time:
Power Saving Idle Time – Parks the read/write heads
Power Saving Standby Time – Lowers disk rotation speed
Power Saving Stopped Time – Spins down the disk (stops rotation)
You must also enable Power Management on the disk array. See "Creating a Disk Array – Advanced (CLU)" on page 133 and "Enabling Media Patrol, PDM, and Power Management on a Disk Array (CLU)" on page 135.

6. Press **Control-A** to save your settings.



Notes

Power Management must be enabled on the disk array for the HDD Power Saving settings to be effective. See "Making Disk Array Settings (CLU)" on page 134

Power management is limited to the features your HDDs actually support.

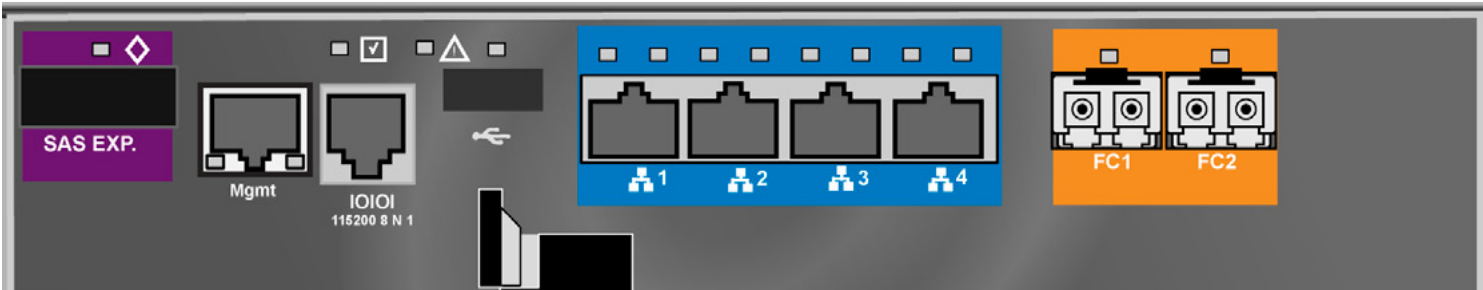
LOCATING THE CONTROLLER (CLU)

To locate this controller:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Controller Management** and press **Enter**.
3. Highlight the controller you want and press **Enter**.
4. Highlight **Controller Settings** and press **Enter**.
5. Highlight **Locate Controller** and press **Enter**.

The controller LEDs blink for one minute.

Figure 5 - 4: Vess R2600fi RAID controller LEDs



MANAGING THE ENCLOSURE (CLU)

Enclosure Management includes the following functions:

- "Viewing the Enclosures Summary (CLU)" on page 126
- "Viewing Enclosure Information (CLU)" on page 126
- "Making Enclosure Settings (CLU)" on page 126
- "Viewing Power Supply Status (CLU)" on page 127
- "Locating a Power Supply (CLU)" on page 127
- "Viewing Cooling Unit Status (CLU)" on page 127
- "Viewing Temperature Sensor Status (CLU)" on page 127
- "Viewing Voltage Sensor Status (CLU)" on page 127
- "Viewing Battery Information (CLU)" on page 128
- "Locating an Enclosure (CLU)" on page 128

VIEWING THE ENCLOSURES SUMMARY (CLU)

Enclosure Management includes information, status, settings and location. To access Enclosure Management:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.

The following information is shown:

- Enclosure ID number
- Enclosure Type
- Operational Status
- Status Description (specific components in need of attention, if any)

VIEWING ENCLOSURE INFORMATION (CLU)

To view enclosure information:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.

You can monitor power supplies, cooling units, enclosure temperatures and voltages, and the battery.

ADJUSTABLE ITEMS

You can set or adjust the following items:

- Enclosure Warning and Critical temperature thresholds
- Controller Warning and Critical temperature thresholds

See "Making Enclosure Settings" below.

For information on Enclosure problems, see "Enclosure Problems" on page 209.

MAKING ENCLOSURE SETTINGS (CLU)

To make Enclosure settings:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Enclosure Settings** and press **Enter**.
5. Highlight the Temperature Warning threshold you want to change.
6. Press the backspace key to erase the current value.
7. Type a new interval value in degrees C.
8. Press **Control-A** to save your settings.

VIEWING FRU VPD INFORMATION (CLU)

FRU VPD refers to Vital Product Data (VPD) information about Field Replaceable Units (FRU) in the enclosure.

The number and type of FRU depends on the subsystem model.

To view FRU VPD information:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **FRU VPD Information** and press **Enter**.

Use this information when communicating with Technical Support and when ordering replacement units.

For contact information, see "Contacting Technical Support" on page 228.

VIEWING POWER SUPPLY STATUS (CLU)

To view the status of the power supplies:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Power Supplies** and press **Enter**.

The screen displays the operational and fan status of Vess R2600’s three power supplies. If any status differs from normal or the fan speed is below the Healthy Threshold value, there is a fan/power supply malfunction. See "Replacing a Power Supply" on page 177.

LOCATING A POWER SUPPLY (CLU)

To locate a power supply:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Power Supplies** and press **Enter**.

5. Highlight **Locate Power Supply** and press **Enter**.

The LED on the selected power supply blinks for one minute.

VIEWING COOLING UNIT STATUS (CLU)

To view the status of the power supply fans:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Cooling Units** and press **Enter**.

The screen displays the status and speed of Vess R2600’s cooling units, which are the power supply fans.

If fan speed is below the Healthy Threshold, there is a malfunction.

VIEWING TEMPERATURE SENSOR STATUS (CLU)

To view the status of the temperature sensors:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Temperature Sensors** and press **Enter**.

If any temperature exceeds the Healthy Threshold value, there is an overheat condition in the enclosure.

See Making Enclosur e Settings (CLU) and see "Diagnosing an Enclosure Problem" on page 209.

VIEWING VOLTAGE SENSOR STATUS (CLU)

To view the status of the voltage sensors:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.

4. Highlight **Voltage Sensors** and press **Enter**.

If any voltage is outside the Healthy Threshold values, there is a voltage malfunction in the enclosure.

See "Diagnosing an Enclosure Problem" on page 209.

VIEWING BATTERY INFORMATION (CLU)

This feature enables you monitor and recondition the subsystem battery or batteries.

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Batteries** and press **Enter**.
5. Highlight the battery you want to monitor and press **Enter**.

BATTERY NOTES

If a battery does not reflect normal conditions and it is not currently under reconditioning, run the Recondition function before you replace the battery. See Reconditioning a Battery (CLU).

Reconditioning fully discharges, then fully recharges the battery. During reconditioning, if the Adaptive Writeback Cache function is enabled, the controller cache is set to **Write Thru**. After reconditioning, the cache is reset to **Write Back**. See "Making Controller Settings (CLU)" on page 124.

If a battery reaches the threshold temperature while charging or discharging, the charge or discharge pauses and the blower runs at high speed until the battery temperature falls below the threshold.

If the battery does not maintain normal values after a Recondition, replace the battery. See "Replacing a Cache Backup Battery" on page 179.

RECONDITIONING A BATTERY (CLU)

To recondition the subsystem battery:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Batteries** and press **Enter**.
5. Highlight the battery you want to recondition and press **Enter**.
6. Highlight **Start Reconditioning** and press **Enter**.
7. Press Y to confirm.

Reconditioning fully discharges, then fully recharges the battery. During reconditioning, if the Adaptive Writeback Cache function is enabled, the controller cache is set to **Write Thru**. After reconditioning, the cache is reset to **Write Back**. See "Making Controller Settings (CLU)" on page 124.

LOCATING AN ENCLOSURE (CLU)

This feature helps you identify the physical Vess R2600 enclosure you are working with through the CLU.

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Management** and press **Enter**.
3. Highlight the enclosure you want and press **Enter**.
4. Highlight **Locate Enclosure** and press **Enter**.

The LEDs on the front of the Vess R2600 blink for one minute.

VIEWING ENCLOSURE TOPOLOGY (CLU)

This feature displays the connection topology of the Vess R2600 subsystem. Topology refers to the manner in which the data paths among the enclosures are connected. There are three methods:

- **Individual Subsystem** – A single subsystem
- **JBOD Expansion** – Managed through one subsystem or head unit
- **RAID Subsystem Cascading** – Managed through one subsystem or head unit

For more information about connections, see "Making Management and Data Connections" on page 18.

To view enclosure topology:

1. From the Main Menu, highlight **Subsystem Management** and press **Enter**.
2. Highlight **Enclosure Topology** and press **Enter**.

The following information applies to the Head Unit:

- **Enclosure number** – 1
- **Controller number** – 1 or 2
- **Port number**
- **Status** – OK is normal. N/C is not connected
- Link Width

The following information applies to RAID cascaded units or JBOD expansion units:

- **Connected EnclWWN** – The subsystem identified by its World Wide Name (WWN)
- **Connected (Encl,Ctrl,Port)** – The subsystem’s enclosure, controller, and port numbers where the data connection was made
- If there is no connection, the value shows N/A.

PHYSICAL DRIVE MANAGEMENT (CLU)

Physical Drive Management includes the following functions:

- "Viewing a List of Physical Drives (CLU)" on page 129
- "Making Global Physical Drive Settings (CLU)" on page 129
- "Viewing Physical Drive Information (CLU)" on page 130
- "Setting an Alias (CLU)" on page 130
- "Clearing Stale and PFA Conditions (CLU)" on page 130
- "Forcing a Physical Drive Offline (CLU)" on page 131
- "Locating a Physical Drive (CLU)" on page 131

VIEWING A LIST OF PHYSICAL DRIVES (CLU)

To view a list of physical drives:

From the Main Menu, highlight **Physical Drive Management** and press **Enter**.

The list of physical drives displays.

MAKING GLOBAL PHYSICAL DRIVE SETTINGS (CLU)

All physical drive settings are made globally, except for setting an alias, which applies to individual drives.

To make global physical drive settings:

1. From the Main Menu, highlight **Physical Drive Management** and press **Enter**.
2. Highlight **Global Physical Drives Settings** and press **Enter**.
3. Change the following settings as required.

For SATA drives:

- Highlight **Write Cache** and press the spacebar to toggle between **Enabled** and **Disabled**.
- Highlight **Read Look Ahead Cache** and press the spacebar to toggle between **Enabled** and **Disabled**.
- Highlight **CmdQueuing** and press the spacebar to toggle between **Enabled** and **Disabled**.
- Highlight **MediumErrorThreshold** and press the backspace key to remove the current value, then type a new smaller value.
See the comments on the next page.

- Highlight **DMA Mode** and press the spacebar to toggle through UDMA 0 to 6 and MDMA 0 to 2.

For SAS drives:

- Highlight **Write Cache** and press the spacebar to toggle between **Enabled** and **Disabled**.
- Highlight **Read Look Ahead Cache** and press the spacebar to toggle between **Enabled** and **Disabled**.
- Highlight **CmdQueuing** and press the spacebar to toggle between **Enabled** and **Disabled**.
- Highlight **MediumErrorThreshold** and press the backspace key to remove the current value, then type a new smaller value.
See the comments below.
- Highlight **Read Cache** and press the spacebar to toggle between **Enabled** and **Disabled**.

4. Press **Control-A** to save your settings.

See “Viewing Physical Drive Information” below to determine which functions your physical drives support.

Medium Error Threshold is the number of bad blocks tolerated before the controller marks the drive as Dead. The default setting is 64 blocks. A setting of zero disables the function. When disabled, no drives are marked offline even when errors are detected.

VIEWING PHYSICAL DRIVE INFORMATION (CLU)

To view information about a physical drive:

1. From the Main Menu, highlight **Physical Drive Management** and press **Enter**.
2. Highlight the physical drive you want and press **Enter**.

Basic information displays.

3. Highlight **Advanced Information** and press **Enter**.

Advanced information displays.

VIEWING PHYSICAL DRIVE STATISTICS (CLU)

To view the statistics for the selected physical drive:

1. From the Main Menu, highlight **Physical Drive Management** and press **Enter**.
2. Highlight the physical drive you want and press **Enter**.
3. Highlight **Physical Drive Statistics** and press **Enter**.

CLEARING STATISTICS

To clear physical drive statistics, see "Clearing Statistics (CLU)" on page 167.

SETTING AN ALIAS (CLU)

An alias is optional. To set an Alias for a physical drive:

1. From the Main Menu, highlight **Physical Drive Management** and press **Enter**.
2. Highlight the physical drive you want and press **Enter**.
3. Type an alias into the field provided.

Maximum of 32 characters. Use letters, numbers, space between words and underscore.

4. Press **Control-A** to save your settings.

CLEARING STALE AND PFA CONDITIONS (CLU)

The Clear Stale and Clear PFA functions only appear when those conditions exist on the physical drive. To clear a Stale or PFA condition on a physical drive:

1. From the Main Menu, highlight **Physical Drive Management** and press **Enter**.
2. Highlight the physical drive you want and press **Enter**.
3. Highlight **Clear Stale** or **Clear PFA** and press **Enter**.

If a physical drive is still online and shows a PFA error but “Clear PFA” does not appear, use PDM to copy the data to a new physical drive. See “Running PDM on a Disk Array” on page 136.

If a physical drive is offline and shows a PFA error, rebuild the disk array. See “Rebuilding a Disk Array”

on page 135. After rebuilding, the drive shows Stale. Run **Clear Stale** then run **Clear PFA**.

If the physical drive with a PFA error is a spare, you must delete the drive as a spare, then **Clear PFA** is available.

After you clear a PFA error, watch for another PFA error to appear. If it does, replace the physical drive.

FORCING A PHYSICAL DRIVE OFFLINE (CLU)

This function enables you to force an online physical drive to go Offline.

The Force Offline function appears only for physical drives that are assigned to disk arrays.



Caution

Forcing a physical drive offline is likely to cause data loss. Back up your data before you proceed. Use this function only when required.



Important

Forcing a physical drive offline causes your logical drives to become degraded. If Auto Rebuild is enabled and a spare drive is available, the disk array begins rebuilding itself automatically.

To force a physical drive offline:

1. From the Main Menu, highlight **Physical Drive Management** and press **Enter**.
2. Highlight **Global Physical Drives Settings** and press **Enter**.
3. Highlight the physical drive you want and press **Enter**.
4. Highlight **Force Offline** and press **Enter**.
5. Press Y to confirm.

LOCATING A PHYSICAL DRIVE (CLU)

This feature helps you identify a physical drive within the Vess R2600 enclosure you are working with through the CLU. To locate a physical drive:

1. From the Main Menu, highlight **Physical Drive Management** and press **Enter**.
2. Highlight **Global Physical Drives Settings** and press **Enter**.
3. Highlight the physical drive you want and press **Enter**.
4. Highlight **Locate Physical Drive** and press **Enter**.

The drive carrier status LED flashes for one minute.

Figure 5 - 5: Drive carrier status LED



MANAGING DISK ARRAYS (CLU)

Disk Array Management includes the following functions:

- Viewing a List of Disk Arrays (page 132)
- Creating a Disk Array (page 132)
- Deleting a Disk Array (page 134)
- Making Disk Array Settings (page 134)
- Viewing Disk Array Information (page 134)
- "Enabling Media Patrol, PDM, and Power Management on a Disk Array (CLU)" on page 135.
- Preparing the Disk Array for Transport (page 135)
- Rebuilding a Disk Array (page 135)
- Running PDM on a Disk Array (page 136)
- Running Transition on a Disk Array (page 136)
- Locating a Disk Array (page 136)

VIEWING A LIST OF DISK ARRAYS (CLU)

To view a list of disk arrays:

From the Main Menu, highlight **Disk Array Management** and press **Enter**.

The list of disk arrays displays.

CREATING A DISK ARRAY (CLU)

The CLU provides three methods of creating a disk array:

- **Automatic** – Creates a new disk array following a default set of parameters. Creates a hot spare drive for all RAID levels except RAID 0, when five or more unconfigured physical drives are available. You can accept or reject the proposed arrangement but you cannot modify it. See “Creating a Disk Array – Automatic”
- **Express** – You choose the parameters for a new disk array by specifying the characteristics you want. You can create multiple logical drives at the same time, however they are all identical. Creates a hot spare drive for all RAID levels except RAID 0. See “Creating a Disk Array – Express”
- **Advanced** – Enables you to specify all parameters for a new disk array, logical drives and spare drives. See

“Creating a Disk Array – Advanced”

CREATING A DISK ARRAY – AUTOMATIC (CLU)

To create a disk array using the Automatic feature:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight **Create New Array** and press **Enter**.
3. Highlight **Configuration Method** and press the spacebar to toggle to **Automatic**.
4. Press **Control-A** to save your settings and move to the next screen.
5. Review the proposed configuration of disk array and logical drives.
 - To accept the proposed configuration and create the disk array and logical drives, highlight **Save Configuration** and press **Enter**.
 - To reject the proposed configuration, highlight **Cancel Array Configuration** and press **Enter**. You return to the Disk Arrays Summary screen.To create a disk array with different characteristics, repeat the steps above specifying different parameters but choose the **Express** or **Advanced** option.

CREATING A DISK ARRAY – EXPRESS (CLU)

To create a disk array using the Express feature:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight **Create New Array** and press **Enter**.
3. Highlight **Configuration Method** and press the spacebar to toggle to **Express**.
4. Highlight the following options and press to spacebar to choose **Yes** or **No**:
 - Redundancy
 - Capacity
 - Performance
 - Spare Drive
 - Mixing SATA/SAS DriveIf you choose No, and you have both SATA and SAS drives, a separate array is created for each type of drive.
5. Highlight **Number of Logical Drives** and press the backspace key to erase the current value, then enter the number of logical drives you want.
6. Highlight **Application Type** and press the spacebar to toggle though the applications and choose the best one for your disk array.

- **File Server**
- **Video Stream**
- **Transaction Data**
- **Transaction Log**
- **Other**

7. Press **Control-A** to save your settings and move to the next screen.
8. Review the proposed configuration of disk array and logical drives.

To accept the proposed configuration and create the disk array and logical drives, highlight **Save Configuration** and press **Enter**.

To reject the proposed configuration, highlight **Cancel Array Configuration** and press **Enter**. You return to the Disk Arrays Summary screen.

To create a disk array with different characteristics, highlight **Create New Array** and press **Enter**. Repeat the steps above specifying different parameters. Or choose the **Advanced** option.

CREATING A DISK ARRAY — ADVANCED (CLU)

For more information on the choices below, see “Chapter 7: Technology Background”

To create a disk array using the Advanced feature:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight **Create New Array** and press **Enter**.
3. Highlight **Configuration Method** and press the spacebar to toggle to **Advanced**.

Step 1 – Disk Array Creation

1. If you want to specify an alias to the disk array, highlight **Alias** and type a name.

Maximum of 32 characters. Use letters, numbers, space between words and underscore.

2. Choose whether to enable Media Patrol, PDM, and Power Management.
3. Choose a Media Type, HDD or SSD.
4. Highlight **Save Settings** and Continue and press **Enter**.
5. Highlight a physical drive you want to add to your array and press the spacebar to choose it.

Repeat this action until you have selected all the physical drives for your array.

6. Highlight **Save Settings and Continue** and press **Enter**.

Step 2 – Logical Drive Creation

1. If you want to specify an alias to the logical drive, highlight **Alias** and type a name.

Maximum of 32 characters. Use letters, numbers, space between words and underscore.

2. Highlight **RAID Level** and press the spacebar to toggle though a list of available RAID levels.
3. If you want to create multiple logical drives, highlight **Capacity**, press the backspace key to remove the current value, then type a new smaller value.
4. RAID 30, 50 and 60 only. Highlight **Number of Axles** and press the spacebar to choose the number of axles.

See "RAID 30 Axles" on Page 187, "RAID 60 Axles" and "RAID 50 Axles" on page 188 .

5. For the following items, accept the default value or highlight and press the spacebar to choose a new value:
 - Highlight **Stripe** and press the spacebar to toggle through stripe sizes and choose 64 KB, 128 KB, 256 KB, 512 KB, or 1 MB.
 - Highlight **Sector** and press the spacebar to toggle through sector sizes and choose 512 B, 1 KB, 2 KB, or 4 KB.
 - Highlight **Write Policy** and press the spacebar to toggle write cache policy between **WriteBack** and **WriteThru** (write though).
 - Highlight **Read Policy** and press the spacebar to toggle read cache policy though **ReadCache**, **ReadAhead**, and **NoCache**.
 - Highlight **Preferred Controller ID** and press the spacebar to toggle among **1**, **2**, or **Automatic**.
 - Highlight PerfectRebuild and press the spacebar to toggle Enable or disable.

6. Highlight **Save Logical Drive** and press **Enter**.

Step 3 – Summary

Review logical drives you are about to create for your new array. Then do one of the following actions:

- If you agree with the logical drives as specified, highlight **Complete Disk Array Creation** and press **Enter**.
- If you specified less than the full capacity for the logical drive in the previous screen, and you want to add another logical drive now, highlight **Create New Logical Drive** and press **Enter**.
- If you do not agree with the logical drives, highlight **Return to Previous Screen** and press **Enter** to begin the process again.

DELETING A DISK ARRAY (CLU)



Caution

When you delete a disk array, you delete all the logical drives and the data they contain. Back up all important data before deleting a disk array.

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want to delete and press the spacebar to mark it.

The mark is an asterisk (*) to the left of the listing.

3. Highlight **Delete Marked Arrays** and press **Enter**.
4. Press Y to confirm the deletion.
5. Press Y again to reconfirm.

MAKING DISK ARRAY SETTINGS (CLU)

To make disk array settings:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.

The list of disk arrays appears.

2. Highlight the disk array you want and press the **Enter**.
3. Make settings changes as required:
 - **Enter**, change or delete the alias in the **Alias** field
Maximum of 32 characters; letters, numbers, space between characters, and underline.
 - **Media Patrol** – Highlight and press the spacebar to toggle between enable and disable.
 - **PDM** – Highlight and press the spacebar to toggle between enable and disable.
 - **Power Management** – Highlight and press the spacebar to toggle between enable and disable.

4. Press **Control-A** to save your settings.



Notes

You can also enable or disable Media Patrol for the entire RAID system. See "Making Background Activity Settings (CLU)" on page 154.

Power Management must be enabled on the disk array for the HDD Power Saving settings to be effective. See Making Disk Array Settings (CLU).

Power management is limited to the features your HDDs actually support.

VIEWING DISK ARRAY INFORMATION (CLU)

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.

The information and settings screen appears.

3. Highlight any of the following and press **Enter** to view a list of:
 - Physical drives in this array
 - Logical drives in this array
 - Spare drives in this array, dedicated and global

DISK ARRAY OPERATIONAL STATUS

- **OK** – This is the normal state of a logical drive. When a logical drive is Functional, it is ready for immediate use. For RAID Levels other than RAID 0 (Striping), the logical drive has full redundancy.
- **Synchronizing** – This condition is temporary. Synchronizing is a maintenance function that verifies the integrity of data and redundancy in the logical drive. When a logical drive is Synchronizing, it functions and your data is available. However, access is slower due to the synchronizing operation.
- **Critical/Degraded** – This condition arises as the result of a physical drive failure. A degraded logical drive still functions and your data is still available. However, the logical drive has lost redundancy (fault tolerance). You must determine the cause of the problem and correct it.
- **Rebuilding** – This condition is temporary. When a physical drive has been replaced, the logical drive automatically begins rebuilding in order to restore redundancy (fault tolerance). When a logical drive is rebuilding, it functions and your data is available. However, access is slower due to the rebuilding operation.
- **Transport Ready** – After you perform a successful Prepare for Transport operation, this condition means you can remove the physical drives of this disk array and move them to another enclosure or different drive slots. After you relocate the physical drives, the disk array status shows OK.

ACCEPTING AN INCOMPLETE ARRAY (CLU)

This condition is the result of a missing physical drive. See “Incomplete Array” on page 215 before you use this function.

To accept an incomplete array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Accept Incomplete Array** and press **Enter**.

ENABLING MEDIA PATROL, PDM, AND POWER MANAGEMENT ON A DISK ARRAY (CLU)

Media Patrol checks the magnetic media on physical drives. Predictive Data Migration (PDM) migrates data from the suspect physical drive to a spare drive **before** the physical drive fails. Power Management parks the heads, spins down, and stops rotation after a set period of time to reduce power consumption.

Media Patrol, PDM, and Power Management are enabled by default. Enabled is the recommended setting for both features.

To enable Media Patrol, PDM, and Power Management on a disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Media Patrol** and press the spacebar to toggle between **Enable** and **Disable**.
4. Highlight **PDM** and press the spacebar to toggle between **Enable** and **Disable**.
5. Highlight **Power Management** and press the spacebar to toggle between **Enable** and **Disable**.
6. Press **Control-A** to save your settings.

See "Running PDM on a Disk Array (CLU)" on page 136 and "Making Background Activity Settings (CLU)" on page 154.

For Power Management settings, see "Making Controller Settings (CLU)" on page 124.

PREPARING THE DISK ARRAY FOR TRANSPORT (CLU)

To run the Transport function on a disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Transport** and press **Enter**.
4. Press Y to confirm.

REBUILDING A DISK ARRAY (CLU)

Before you can rebuild, you must have a replacement or target physical drive of adequate capacity for your disk array.

To rebuild a disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Background Activities** and press **Enter**.
4. Highlight **Rebuild** and press **Enter**.

Default source and target drives are shown with possible alternative choices.

5. To choose different drive, highlight the drive, press the backspace key to remove the current number, then type a new number.
6. Highlight **Start** and press **Enter**.

For rebuild rate, see "Making Background Activity Settings (CLU)" on page 154.

RUNNING MEDIA PATROL ON A DISK ARRAY (CLU)

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. If Media Patrol encounters a critical error, it triggers PDM if PDM is enabled on the disk array.

See "Enabling Media Patrol, PDM, and Power Management on a Disk Array (CLU)" on page 135.

For Media Patrol rate, see "Making Background Activity Settings (CLU)" on page 154.

RUNNING PDM ON A DISK ARRAY (CLU)

Predictive Data Migration (PDM) migrates data from the suspect physical drive to a spare drive **before** the physical drive fails.

Before you can run PDM, you must have a replacement or target physical drive of adequate capacity for your disk array.

To run PDM on a disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Background Activities** and press **Enter**.
4. Highlight **Predictive Data Migration** and press **Enter**.

Default source and target drives are shown with possible alternative choices.

5. To choose different drive, highlight the drive, press the backspace key to remove the current number, then type a new number.
6. Highlight **Start** and press **Enter**.

See "Enabling Media Patrol, PDM, and Power Management on a Disk Array (CLU)" on page 135.

For PDM rate, see "Making Background Activity Settings (CLU)" on page 154.

RUNNING TRANSITION ON A DISK ARRAY (CLU)

Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare drive. For more information, see “Transition” on page 193.

In order to run Transition:

- The spare drive must be Revertible.
- You must have an unconfigured physical drive of the same or larger capacity to replace the spare drive.

To run Transition on a disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Background Activities** and press **Enter**.

4. Highlight **Transition** and press **Enter**.

Default source and target drives are shown with possible alternative choices.

5. To choose different drive, highlight the drive, press the backspace key to remove the current number, then type a new number.
6. Highlight **Start** and press **Enter**.

For transition rate, see "Making Background Activity Settings (CLU)" on page 154.

LOCATING A DISK ARRAY (CLU)

This feature helps you identify the physical drives assigned to the disk array you are working with in the CLU.

To locate a disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Locate Disk Array** and press **Enter**.

The drive carrier status LEDs flash for one minute.

MANAGING SPARE DRIVES (CLU)

Spare Drive Management includes the following functions:

- "Viewing a list of Spare Drives (CLU)" on page 137
- "Creating a Spare Drive (CLU)" on page 137
- "Making Spare Drive Settings (CLU)" on page 137
- "Running Spare Check (CLU)" on page 138
- "Deleting a Spare Drive (CLU)" on page 138

VIEWING A LIST OF SPARE DRIVES (CLU)

To view a list of spare drives:

From the Main Menu, highlight **Spare Drive Management** and press **Enter**.

A list of the current spare drives appears, including the following parameters:

- **ID number**
- **Operational Status**
- **Physical Drive ID number**
- **Configured Capacity**
- **Revertible** – The spare drive returns to spare status after you replace the failed drive in the disk array. See “Transition” on page 356 for more information.
- **Type** – Global (all disk arrays) or Dedicated (to specified disk arrays)
- **Dedicated to Array** – The array to which a dedicated spare is assigned

For more information, see "Disk Arrays" on page 183.

CREATING A SPARE DRIVE (CLU)

Only unconfigured physical drives can be used to make spares. Check your available drives under Physical Drive Management. Also see "Managing Physical Drives" on page 87.

1. From the Main Menu, highlight **Spare Drive Management** and press **Enter**.
2. Highlight **Create New Spare Drive** and press **Enter**.

A default physical drive is shown with possible alternative choices.

3. To choose different drive, highlight the drive, press the backspace key to remove the current number, then

type a new number.

4. Highlight **Revertible** and press the spacebar to toggle between **Yes** and **No**.

A revertible drive can be returned to spare status after you replace the failed drive in a disk array. See "Transition" on page 193 for more information.

5. Highlight **Spare Type** and press the spacebar to toggle between **Dedicated** and **Global**.

Dedicated means this spare drive can only be used with the specified disk arrays. Global means this spare drive can be used by any disk array.

If you chose Dedicated, a default disk array is shown with possible alternative choices.

To choose different array, highlight the array and press the backspace key to erase the current number, then type the new number.

6. Press **Control-A** to save the spare drive.

MAKING SPARE DRIVE SETTINGS (CLU)

To change spare drive settings:

1. From the Main Menu, highlight **Spare Drive Management** and press **Enter**.

A list of the current spare drives appears, including the following parameters:

2. Highlight the spare drive you want to change and press **Enter**.
3. Highlight the setting you want to change:
 - **Revertible** – A revertible drive can be returned to spare status after you replace the failed drive in a disk array. See “Transition” on page 356 for more information.
 - **Type** – Dedicated means this spare drive can only be used with the specified disk arrays. Global means this spare drive can be used by any disk array.

4. Press the spacebar to toggle between the choices.
5. For dedicated spares, type the array number the spare is assigned to.
6. Press **Control-A** to save your settings.

RUNNING SPARE CHECK (CLU)

To run Spare Check:

1. From the Main Menu, highlight **Spare Drive Management** and press **Enter**.

A list of the current spare drives appears.

2. Highlight the spare drive you want to check and press **Enter**.
3. Highlight **Start Spare Check** and press **Enter**.

The results appear next to Spare Check Status in the same window. Healthy means normal.

DELETING A SPARE DRIVE (CLU)



Caution

If the spare drive you delete is the only spare, the controller does not rebuild a critical array until you provide a new spare drive.

To delete a spare drive:

1. From the Main Menu, highlight **Spare Drive Management** and press **Enter**.

A list of the current spare drives appears.

2. Highlight the spare drive you want to delete and press the spacebar to mark it.

The mark is an asterisk (*) to the left of the listing.

3. Highlight **Delete Marked Spare Drives** and press **Enter**.
4. Press Y to confirm the deletion.

MANAGING LOGICAL DRIVES (CLU)

Logical drive management includes:

- "Creating a Logical Drive (CLU)" on page 138
- "Deleting a Logical Drive (CLU)" on page 139
- "Viewing the Logical Drive Check Table (CLU)" on page 139
- "Making Logical Drive Settings (CLU)" on page 140
- "Initializing a Logical Drive (CLU)" on page 140
- "Running Redundancy Check (CLU)" on page 140
- "Locating a Logical Drive (CLU)" on page 140
- "Migrating a Logical Drive (CLU)" on page 141
- "Creating a LUN Clone (CLU)" on page 141

For LUN mapping, see "Working with LUN Mapping (CLU)" on page 156

CREATING A LOGICAL DRIVE (CLU)

You can create logical drives on existing disk arrays if there is available space in the array.

To create a logical drive from an existing disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array in which you want to create a logical drive and press **Enter**.
3. Highlight **Logical Drives in the Disk Array** and press **Enter**.
4. Highlight **Create New Logical Drive** and press **Enter**.

The Disk Array ID number and Maximum capacity available for the new logical drive are displayed.

5. Highlight the following parameters and press the backspace key to erase the current value:
 - **Alias** – Type an alias into the field, if desired. Maximum of 32 characters. Use letters, numbers, space between words and underscore.
 - **RAID Level** - Press the spacebar to toggle though a list of available RAID levels.
 - **Capacity** – Maximum capacity shown. **Enter** a smaller capacity if desired.

6. Highlight the following parameters and press the spacebar to toggle though the available choices:
- **Stripe size** – Press the spacebar to choose: 64 KB, 128 KB, 256 KB, 512 KB, or 1 MB.
 - **Sector size** – Press the spacebar to choose: 512 B; 1 KB, 2 KB, or 4 KB.
 - **Write Policy** – Press spacebar to choose: Write Back or Write Through.
 - **Read Policy** – Press spacebar to choose: No Cache, Read Cache, or Read Ahead Cache.
7. Highlight **Preferred Controller ID** and press the spacebar to toggle among **1**, **2**, or **Automatic**.
8. Highlight PerfectRebuild and press the spacebar to toggle Enable or disable.
9. RAID 30, 50 and 60 only. Highlight **Number of Axles** and press the spacebar to choose the number of axles.
- 10.Highlight **Save Logical Drive** and press **Enter**.



Note

If you did not use all of the available capacity of the disk array, you can create an additional logical drive at this point.

DELETING A LOGICAL DRIVE (CLU)



Caution

When you delete a logical drive, you delete all the data it contains. Back up all important data before deleting a logical drive.

To delete a logical drive from a disk array:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array that contains the logical drive you want to delete and press **Enter**.
3. Highlight **Logical Drives in the Disk Array** and press **Enter**.
4. Highlight the logical drive you want to delete and press the spacebar to mark it.

The mark is an asterisk (*) to the left of the listing.

5. Highlight **Delete Marked Logical Drives** and press **Enter**.
6. Press Y to confirm the deletion.

Press Y again to re-confirm.

VIEWING LOGICAL DRIVE INFORMATION (CLU)

To view logical drive information:

1. From the Main Menu, highlight **Logical Drive Management** and press **Enter**.
2. Highlight the logical drive you want and press **Enter**.

The information and settings screen appears.

3. Highlight any of the following and press **Enter** to view more information:
 - **Check Table** – Read Check, Write Check, and Inconsistency Check Tables
 - **Logical Drive Statistics**

VIEWING LOGICAL DRIVE STATISTICS (CLU)

To view logical drive information:

1. From the Main Menu, highlight **Logical Drive Management** and press **Enter**.
2. Highlight the logical drive you want and press **Enter**.

The information and settings screen appears.

3. Highlight **Logical Drive Statistics** and press **Enter**.

The statistics screen appears.

To clear logical drive statistics, see "Clearing Statistics (CLU)" on page 167.

VIEWING THE LOGICAL DRIVE CHECK TABLE (CLU)

To view logical drive information:

1. From the Main Menu, highlight Logical Drive Management and press **Enter**.
2. Highlight the logical drive you want and press **Enter**.
3. Highlight Check Table and press **Enter**.
4. Highlight one of the following options and press **Enter**:

- **Show All Records**
- **Read Check Table**
- **Write Check Table**
- **Inconsistent Check Table**

MAKING LOGICAL DRIVE SETTINGS (CLU)

To make Logical Drive settings:

1. From the Main Menu, highlight **Logical Drive Management** and press **Enter**.
2. Highlight the logical drive you want and press **Enter**.
3. For the following items, accept the existing setting choose a new one:
 - Highlight **Alias** and type an alias into the field provided.
 - Maximum of 32 characters. Use letters, numbers, space between words and underscore. An alias is optional.
 - Highlight **WritePolicy** and press the spacebar to toggle between **WriteBack** and **WriteThru** (write though).
 - Highlight **ReadPolicy** and press the spacebar to toggle though **ReadCache**, **ReadAhead** and **None**.
 - Highlight **Preferred Controller ID** and press the spacebar to toggle between **1** and **2**.
 - Highlight **PerfectRebuild** and press the spacebar to toggle between Enable and Disable. Note that once PerfectRebuild is disbaled it can not be enabled again.
4. Press **Control-A** to save your settings.

INITIALIZING A LOGICAL DRIVE (CLU)

This function sets all data bits in the logical drive to zero.



Warning

When you initialize a logical drive, all the data on the logical drive is lost. Backup any important data before you initialize a logical drive.

To initialize a logical drive:

1. From the Main Menu, highlight **Logical Drive Management** and press **Enter**.
2. Highlight the logical drive you want and press **Enter**.
3. Highlight **Background Activities** and press **Enter**.
4. Highlight **Start Initialization** and press **Enter**.

The initialization parameters appear.

- **Initialization pattern** – The default 00000000 is best for most applications
- **Quick Initialization** – Yes means only the first and last sections of the logical drives are initialized. No means the entire logical drive is initialized.

To change a parameter, highlight it and press the backspace key to erase the current value, then type the

new value.

5. Highlight **Start** and press **Enter**.

If necessary, you can pause and resume or stop and restart the Initialization. You cannot access the logical drive until Initialization has finished.

For initialization rate, see "Making Background Activity Settings (CLU)" on page 154.

RUNNING REDUNDANCY CHECK (CLU)

Redundancy Check is a maintenance procedure for logical drives in fault-tolerant disk arrays that ensures all the data matches exactly.

To run Redundancy Check:

1. From the Main Menu, highlight **Logical Drive Management** and press **Enter**.
2. Highlight the logical drive you want and press **Enter**.
3. Highlight **Background Activities** and press **Enter**.
4. Highlight **Start Redundancy Check** and press **Enter**.

The redundancy check parameters appear.

- **Auto Fix** – Corrects inconsistencies automatically
- **Pause On Error** – Pauses the Redundancy Check when an error is found

To change a parameter, highlight it and press the backspace toggle between **Yes** and **No**.

5. Highlight **Start** and press **Enter**.

If necessary, you can pause and resume or stop and restart the Redundancy Check. You can use the logical drive while Redundancy Check is running.

For Redundancy Check rate, see "Making Background Activity Settings (CLU)" on page 154.

LOCATING A LOGICAL DRIVE (CLU)

This feature helps you identify the physical drives assigned to the logical drive you are working with in the CLU.

To locate a logical drive:

1. From the Main Menu, highlight **Logical Drive Management** and press **Enter**.
2. Highlight the logical drive you want and press **Enter**.

3. Highlight **Locate Logical Drive** and press **Enter**.

The drive carrier status LEDs flash for one minute.

Figure 5 - 6: Drive carrier status LED



MIGRATING A LOGICAL DRIVE (CLU)

In order to migrate RAID level, you may have to add physical drives. For more information, see "RAID Level Migration" on page 189.

To migrate a logical drive:

1. From the Main Menu, highlight **Disk Array Management** and press **Enter**.
2. Highlight the disk array you want and press **Enter**.
3. Highlight **Background Activities** and press **Enter**.
4. Highlight **Migration** and press **Enter**.
5. Highlight the physical drives you want to add and press the spacebar to choose them.



Note

You can add physical drives to a RAID 50 or 60 array but you cannot change the number of axes.
If you add an odd number of physical drives to a RAID 10 array, it becomes a RAID 1E array by default.

6. Highlight **Save Settings and Continue** and press **Enter**.
7. Highlight a logical drive in the list that you want to migrate and press **Enter**.
8. Highlight **RAID Level** and press the spacebar to toggle through the available RAID levels.
9. Optional. If you want to increase capacity of the logical drive, highlight **Expand Capacity** and press the spacebar to toggle to **Yes**.

Highlight **Capacity**, press the backspace key to erase the current capacity and type in the new value.

The new value must be equal or larger than the current capacity.

10.Highlight **Save Logical Drive** and press **Enter**.

The screen returns to Disk Array Migration Logical Drives.

At this point, if you have other logical drives in the same disk array, you can choose them for migration at the same time.

11.Highlight **Complete Disk Array Migration** and press **Enter**.

12.Press Y to confirm.

The screen returns to Disk Arrays Summary.

For migration rate, see "Making Background Activity Settings (CLU)" on page 154.

CREATING A LUN CLONE (CLU)

A LUN clone is an exact copy of the original LUN or logical drive, including all the data it contains, at one point in time. Use a LUN clone as a backup or to migrate a LUN from one system to another.



Important

The action of creating a LUN momentarily takes the original LUN or logical drive offline, meaning nobody can read or write to it.

A LUN clone has the same capacity, stripe size, read and write policies as the original LUN. However, the LUN clone can be a different RAID level. The choice of RAID levels depends on the disk array. And if you have multiple disk arrays, you can create the LUN clone on a different disk array than the original LUN.

This action requires **Super User** or **Power User** privileges.

To create a LUN clone of a logical drive:

1. From the Main Menu, highlight **Logical Drive Management** and press **Enter**.
2. Highlight the logical drive you want to clone and press **Enter**.
3. Highlight **LUN Clone** and press **Enter**.
4. Highlight the RAID Level of Copies field, type the RAID level you want.
5. Highlight **Save Settings and Continue** and press **Enter**.
6. Highlight the disk array you want to use and press the Spacebar to mark it.
7. Highlight **Save Settings and Continue** and press **Enter**.
8. Highlight the Number of Copies field and type the number of LUN clones you want to create.

You can create up to 8 clones of a LUN at a time.

- 9. Highlight Start and press enter to begin the cloning process.
- 10. Press any key to continue.
- 11. Press Y to confirm LUN clone creation.

The cloning progress bar displays.

Note the **Target Logical Drive ID**. Use this number to identify the LUN clone in the Logical Drive list.

If you chose a redundant RAID level, the LUN clone is automatically synchronized after creation.

After the LUN clone is created, you can manage it like any other logical drive. See Making Spare Drive Settings (CLU), Locating a Logical Drive (CLU), and Deleting a Logical Drive (CLU).

For users to access the LUN clone, you must map it to an initiator. See "Working with LUN Mapping (CLU)" on page 156

MANAGING THE NETWORK CONNECTION (CLU)

Network Management deals with network connections and settings for the Vess R2600's Management ports.

Each Management Port can be configured:

- "Making Virtual Management Port Settings (CLU)" on page 142
- "Making Maintenance Mode Settings (CLU)" on page 143

MAKING VIRTUAL MANAGEMENT PORT SETTINGS (CLU)

The Vess R2600 subsystem has a virtual management port, enabling you to log into a Vess R2600 with dual controllers using one IP address.

Before you change settings, please see "About IP Addresses" on page 36.

You initially made these settings during subsystem setup. You can change them later as required.



Caution

Changing virtual management port settings can interrupt your network connection and require you to log in again.

MAKING AUTOMATIC SETTINGS

Automatic settings require a DHCP server on your network. DHCP is currently supported on IPv4 only.

To enable automatic management port settings:

1. From the Main Menu, highlight **Network Management** and press **Enter**.
2. Highlight the protocol family (IPv4 or IPv6) you want and press **Enter**.
3. Highlight **Network Settings** and press **Enter**.
4. Highlight **DHCP** and press the spacebar to toggle to **Enabled**.
5. Press **Control-A** to save your settings.

MAKING MANUAL SETTINGS

1. From the Main Menu, highlight **Network Management** and press **Enter**.
2. Highlight the protocol family (IPv4 or IPv6) you want and press **Enter**.
3. Highlight **Network Settings** and press **Enter**
4. Highlight **DHCP** and press the spacebar to toggle to **Disabled**.

DHCP is currently supported by and does not appear under IPv6.

5. Highlight each of the following and press the backspace key to erase the current value, then type the new value.
- **IP Address**
 - **Subnet Mask**
 - **Default Gateway IP Address**
 - **DNS Server IP Address**

6. Press **Control-A** to save your settings.

MAKING MAINTENANCE MODE SETTINGS (CLU)

Each controller has its own IP addresses for access when the controller goes into maintenance mode. For more information, see "Maintenance Mode" on page 211.

Before you change settings, please see “About IP Addresses” on page 36.

MAKING AUTOMATIC SETTINGS

1. From the Main Menu, highlight **Network Management** and press **Enter**.
2. Highlight **Maintenance Mode Network Configuration** and press **Enter**.
3. Highlight the controller (CId 1 or 2) and protocol family (IPv4 or IPv6) you want and press **Enter**.
4. Highlight **DHCP** and press the spacebar to toggle to **Enabled**.
5. Press **Control-A** to save your settings.

MAKING MANUAL SETTINGS

1. From the Main Menu, highlight **Network Management** and press **Enter**.
2. Highlight **Maintenance Mode Network Configuration** and press **Enter**.
3. Highlight the controller (CId 1 or 2) and protocol family (IPv4 or IPv6) you want and press **Enter**.
4. Highlight **DHCP** and press the spacebar to toggle to **Disabled**.
5. Highlight each of the following and press the backspace key to erase the current value, then type the new value.
 - **IP Address**
 - **Subnet Mask**
 - **Default Gateway IP Address**
 - **DNS Server IP Address**

6. Press **Control-A** to save your settings.

MANAGING FIBRE CHANNEL CONNECTIONS (CLU)

The Fibre Channel Management option appears only with Vess R2600 Fibre Channel models. Fibre Channel Management includes the following functions:

- "Viewing Node Information (CLU)" on page 143
- "Viewing Fibre Channel Port Information" on page 143
- "Making Fibre Channel Port Settings (CLU)" on page 144
- "Viewing Fibre Channel Port Statistics (CLU)" on page 144
- "Viewing SFP Information (CLU)" on page 144
- "Viewing Fibre Channel Initiators (CLU)" on page 145

Also see: "Adding an FC Initiator" on page 104 and "Deleting an FC Initiator" on page 104

VIEWING NODE INFORMATION (CLU)

These functions affect both Vess R2600 Fibre Channel ports.

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Node** and press **Enter**.

Node information appears. There are no user settings on this screen.

VIEWING FIBRE CHANNEL PORT INFORMATION

To view Fibre Channel port information:

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Ports** and press **Enter**.

Highlight the port you want and press **Enter**.

VIEWING FIBRE CHANNEL LOGGED-IN DEVICES (CLU)

To view a list of logged-in devices:

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Ports** and press **Enter**.
3. Highlight the port you want and press **Enter**.
4. Highlight **Logged In Devices** and press **Enter**.

If a Fibre Channel switch is attached, it also appears in this list.

MAKING FIBRE CHANNEL PORT SETTINGS (CLU)

To make Fibre Channel port settings:

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Ports** and press **Enter**.
3. Highlight the port you want and press **Enter**.
4. Highlight **Fibre Channel Port Settings** and press **Enter**.
5. Highlight the following parameters and press the spacebar to toggle though the choices:
 - **Configured Link Speed** – 8 Gb/s, 4 Gb/s, 2 Gb/s, or Automatic selection
 - **Configured Topology** – NL-Port (Arbitrated Loop), N-Port (Point to Point) or Automatic selection
6. Highlight **Hard ALPA** and press the backspace key to erase the current value, then type the new value.

The range is 0 to 255. 255 disables this feature.

7. Press **Control-A** to save your settings.

The table below shows the type of attached topology you achieve based on your connection type and the configured topology you choose:

Fibre Channel Attached Topology		
	Configured Topology	
Connection Type	N-Port	NL-Port
Switch	Fabric Direct	Public Loop
Direct	Point-to-Point	Private Loop

Example 1: If you connect the Vess R2600 to a Fibre Channel switch and choose NL-Port topology, you create a Public Loop attached topology.

Example 2: If you have a Point to Point attached topology, you made a direct connection (no switch) and chose N-port topology.



Note

In some cases, HBA settings to N-Port only work if connected to the switch. Refer to your HBA manual for more information.

VIEWING FIBRE CHANNEL PORT STATISTICS (CLU)

To view Fibre Channel port statistics:

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Ports** and press **Enter**.
- Highlight the port you want and press **Enter**.
3. Highlight **Fibre Channel Port Statistics** and press **Enter**.

VIEWING SFP INFORMATION (CLU)

To view information about the SFPs (small form-factor pluggable transceivers):

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Ports** and press **Enter**.
3. Highlight the port you want and press **Enter**.
4. Highlight **Fibre Channel Port SFP** and press **Enter**.

The screen displays information about the SFP transceiver. There are no user settings on this screen.

VIEWING FIBRE CHANNEL PORT STATISTICS (CLU)

To view port statistics:

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Ports** and press **Enter**.
3. Highlight the port you want and press **Enter**.
4. Highlight **Fibre Channel Port Statistics** and press **Enter**.

This screen displays statistics for this port. There are no user settings on this screen.

CLEARING STATISTICS

To clear Fibre Channel statistics, see "Clearing Statistics (CLU)" on page 167.

PROPERTY DEFINITIONS

Definitions of the properties for which statistical information is reported appears in the list below.

- **TimeLastReset** – Time in minutes since the system has been running.
- **FramesSent** – Number of frames sent since last reset.
- **FramesReceived** – Number of frames received since last reset.
- **WordsSent** – Number of words sent since last reset.
- **WordsReceived** – Number of words received since last reset.
- **LIPCount** – Loop Initialization Primitive Sequence. This primitive sequence applies only to the arbitrated loop topology. It is transmitted by an L_Port to initialize or re-initialize the loop.
- **NOSCount** – Not Operational Primitive Sequence. This primitive sequence is used during link initialization between two N_Ports in the point-to-point topology or an N_Port and an F_Port in the fabric topology.

NOS is sent to indicate that the transmitting port has detected a link failure or is offline. The expected response to a port sending NOS is the OLS primitive sequence.
- **ErrorFrames** – FC devices propagate handshake signals back-and-forth requesting and acknowledging each byte transferred. FC transfers occur in one frame of data at a time. In this case, the value reflects the number of frames with errors.
- **DumpedFrames** – This field specifies the number of frames dumped due to a lack of host buffers.
- **LinkFailureCount** – Number of times the link has failed. Can be caused by a disconnected link or a bad fiber element.

- **LossSyncCount** – Number of times a loss of sync has occurred since last reset.
- **PrimitiveSeqErrorCount** – An ordered set transmitted repeatedly and used to establish and maintain a link.

LR, LRR, NOS, and OLS are primitive sequences used to establish an active link in a connection between two N_Ports or an N_Port and an F_Port.

LIP, LPB, and LPE are primitive sequences used in the Arbitrated Loop topology for initializing the loop and enabling or disabling an L_Port.
- **InvalidWordSentCount** – Number of invalid words sent since last reset.
- **InvalidCRCCount** – Invalid Cyclical Redundancy Count. Number of frames received with an invalid CRC since last reset.
- **InitiatorIOCount** – I/O Count on the initiator on the host side.

CLEARING STATISTICS

To clear statistics, see "Clearing Statistics (CLU)" on page 167.

VIEWING FIBRE CHANNEL INITIATORS (CLU)

LUN Mapping must be enabled in order for Vess R2600 to recognize a Fibre Channel. See "Enabling LUN Mapping (CLU)" on page 156

To view Fibre Channel initiators:

1. From the Main Menu, highlight **Fibre Channel Management** and press **Enter**.
2. Highlight **Fibre Channel Initiators** and press **Enter**.

A list of all currently logged-in initiators appears on the screen.

MANAGING iSCSI CONNECTIONS (CLU)

- "Viewing iSCSI Target Information (CLU)" on page 147
- "Making iSCSI Target Settings (CLU)" on page 147
- "Viewing a List of iSCSI Ports (CLU)" on page 148
- "Viewing iSCSI Port Information (CLU)" on page 148
- "Viewing a List of iSCSI Portals (CLU)" on page 148
- "Viewing iSCSI Portal Information (CLU)" on page 149
- "Adding iSCSI Portals (CLU)" on page 149
- "Making iSCSI Portal Settings (CLU)" on page 149
- "Deleting iSCSI Portals (CLU)" on page 150
- "Deleting an iSCSI Session (CLU)" on page 150
- "Viewing iSCSI Session Information (CLU)" on page 151
- "Making iSCSI iSNS Settings (CLU)" on page 151
- "Viewing a List of iSCSI CHAPs (CLU)" on page 152
- "Adding iSCSI CHAPs (CLU)" on page 152
- "Deleting iSCSI CHAPs (CLU)" on page 152
- "Pinging a Host or Server on the iSCSI Network (CLU)" on page 152
- "Viewing a List of iSCSI Trunks (CLU)" on page 153
- "Adding iSCSI Trunks (CLU)" on page 153
- "Making iSCSI Trunk Settings (CLU)" on page 153
- "Deleting iSCSI Trunks (CLU)" on page 153

VIEWING A LIST OF iSCSI TARGETS (CLU)

A **target** is a logical drive on the Vess R2600 subsystem.

The default target exposes all logical drives and is associated with all portals on the subsystem.

To view a list of iSCSI targets:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Targets** and press **Enter**.

- The list of iSCSI Targets displays.
- **Id** – Target number. 0 is the default target.
 - **Alias** – User assigned name of the target
 - **AssignedPortals** – portals assigned to the target

VIEWING iSCSI TARGET INFORMATION (CLU)

To view information for an iSCSI target:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight iSCSI Targets and press **Enter**.

The list of **iSCSI Targets** displays.

3. Highlight the target you want to change and press **Enter**.

The target information screen displays. Information includes:

- **TargetName** – iSCSI qualified name (iqn) of this target.
- **TargetAlias** – Maximum of 32 characters. Use letters, numbers, space between words, and underscore. An alias is optional.*
- **TargetStatus** – Up or down.
- **ErrorRecovLevel** – Error recovery level supported.
- **ImmediateData** – Enables the initiator to send unsolicited data with the iSCSI command PDU.
- **MaxConnection** – Maximum number of connections.
- **DataPDUInOrder** – Enables placement of data in PDU order.
- **InitialR2T** – Allows initiator to begin sending data to a target without receiving a ready to transfer command.
- **DataSeqInOrder** – Enables placement of data in sequential order.
- **OutStandingR2T** – Maximum number of R2T PDUs the target can have outstanding for a single iSCSI command.
- **MaxBurstLen** – Maximum length of a solicited data sequence in bytes.
- **DefTimeToWait** – After a dropped connection, the number of seconds to wait before attempting to reconnect.
- **DefTimeToRetain** – Number of seconds after time to wait (above) before reassigning outstanding commands.
- **HeaderDigest** – Enables the use of header digest (CRC). Enabled or disabled.*
- **DataDigest** – Enables the use of a data digest (CRC). Enabled or disabled.*
- **UniCHAPAuthen** – Uni-directional (peer) CHAP authentication, enabled or disabled.*
- **BiCHAPAuthen** – Bi-directional (local) CHAP authentication, enabled or disabled.*
- **FirstBurstLen** – First burst length in bytes.
- **AssignedPortals** – Portals assigned to this target.
- **NOP-In**: Check iSCSI connection status*

Items marked with an asterisk (*) are adjustable under "Making iSCSI Target Settings (CLU)" on page

MAKING iSCSI TARGET SETTINGS (CLU)

To make target settings:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Targets** and press **Enter**.

The list of **iSCSI Targets** displays.

3. Highlight the target you want to change and press **Enter**.

The target information screen displays.

4. Highlight **iSCSI Target Settings** and press **Enter**.
5. Make new settings as needed.
 - Optional. Highlight **TargetAlias** and type an alias into the field provided.
 - Highlight each item and press the Spacebar to toggle between Enable and Disable.
 - **HeaderDigest** – Adds a header digest (CRC).
 - **DataDigest** – Adds a data digest (CRC).
 - **UniCHAPAuthen** – Enables uni-directional CHAP authentication.
 - **BiCHAPAuthen** – Enables bi-directional CHAP authentication. Authentication requires a pre-existing CHAP.
 - **Enable NOP-In** - Enable to check iSCSI connection status
6. Highlight **Save Settings** and press **Enter**.
7. Press **Return to Previous Menu** to return to the iSCSI targets list.

VIEWING A LIST OF iSCSI PORTS (CLU)

An iSCSI port is the physical iSCSI connection on the Vess R2600. There are four iSCSI ports on each RAID controller for a total of eight per subsystem.

To view a list of iSCSI ports:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Ports** and press **Enter**.

The list of ports appears with controller and port numbers.

VIEWING iSCSI PORT INFORMATION (CLU)

To view information for an iSCSI target port:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Ports** and press **Enter**.

The list of ports appears with controller and port numbers.

3. Highlight the port you want to see and press **Enter**.

The target port information screen displays. Information includes:

- **CtrlId** – Controller ID (1 or 2)
- **JumboFrame** – Jumbo frames, enabled or disabled*
- **LinkStatus** – Link status, up or down, Active or Inactive
- **MACAddress** – MAC address of the target port
- **MaxSupportedSpeed** – Maximum speed supported (1 Gb/s)
- **CurrentSpeed** – Current or actual speed of the target port
- **RelativePortals** – The portals corresponding to this target port

Items marked with an asterisk (*) are adjustable under “Making iSCSI Port Settings” below.

MAKING iSCSI PORT SETTINGS (CLU)

To make port settings:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Ports** and press **Enter**.

The list of ports appears with controller and port numbers.

3. Highlight the port you want to change and press **Enter**.

The target port information screen displays.

4. Highlight **iSCSI Port Settings** and press **Enter**.
5. Highlight each item and press the Spacebar to toggle between Enable and Disable as needed.
 - **JumboFrame** – Enables and disables jumbo frame support
6. Highlight **Save Settings** and press **Enter**.
7. Press Y to acknowledge possible interruption of iSCSI services.
8. Press Y again to confirm the changes.
9. Highlight **Return to Previous Menu** and press **Enter** to return to the target port information screen.

VIEWING A LIST OF iSCSI PORTALS (CLU)

A **portal** is the interface between an iSCSI port and the iSCSI network.

To view a list of iSCSI portals:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Portals** and press **Enter**.

The list of iSCSI Portals displays.

- **PortalId** – Portal number. Starts at 0.
- **CtrlId** – RAID controller ID, 1 or 2.
- **PortId** – Physical port on the RAID controller, 1 to 4.
- **TrunkId** – Trunk ID, 1 to 8. Refers to portals associated with a trunk (link aggregation). N/A means this portal is not associated with a trunk.
- **VlanTag** – VLAN Tag, 1 to 4094. Refers to portals associated with a Virtual Local Area Network (VLAN). N/A means this portal is not associated with a VLAN.
- **IP** – IP address of the portal.

VIEWING iSCSI PORTAL INFORMATION (CLU)

To view information for an iSCSI target port:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Portals** and press **Enter**.

The list of portals appears.

3. Highlight the port you want to see and press **Enter**.

The portal information screen displays. Information includes:

- **PortalID** – Portal number. Starts at 0.
- **TcpPort** – TCP port number. 3260 is the default and recommended number.
- **DHCP** – Enabled or disabled.*
DHCP is currently supported only for IPv4.
- **AssociatedType** – PHY, VLAN, or Trunk.
- **ControllerID** – RAID controller ID, 1 or 2.
- **PortID** – Physical port on the RAID controller, 1 to 4.
- **InterfaceName** – device name.
- **ProtocolFamily** – IPv4 or IPv6.*
- **PrimaryIP** – Primary IP address of this portal.*
- **Gateway** - Gateway IP address of this portal*
- **PrimaryIPMask** – Subnet mask of this portal.*
- **AssignedTarget** – 0 is the default target. The number of targets available depends on how many targets you create. See Adding iSCSI Targets (CLU) below.

Items marked with an asterisk (*) are adjustable under "Making iSCSI Portal Settings (CLU)" on page 149.

ADDING iSCSI PORTALS (CLU)

Vess R2600 supports up to 32 iSCSI portals. Each iSCSI portal can belong to a different VLAN for a maximum of 32 VLANs.

If you plan to associate the new portal with a trunk, create the trunk first. See "Adding iSCSI Trunks (CLU)" on page 153.

For more information about iSCSI VLANs, see "iSCSI on a VLAN" on page 197.

To add an iSCSI portal:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.

2. Highlight **iSCSI Portals** and press **Enter**.

The list of iSCSI Portals displays.

3. Highlight **Create New Portal** and press **Enter**.

4. Highlight **AssociatedType** and press the Spacebar to toggle through Physical, VLAN, and Trunk.

5. If you chose:
 - **PHY** – Choose a Controller ID (1 or 2) and a Port ID (1 to 4).
 - **Trunk** – Choose a Trunk ID (1 to 8).

To change an ID number, highlight the item, press Backspace to delete the current ID and type a new ID.

6. If you use Associated PHY, highlight **EnableVLAN** and press the Spacebar to toggle between **Enable** and **Disable**. If you choose **Enable**, enter a Vlan Tag (1 to 4094).

To change a value, highlight the item, press Backspace to delete the current value and type a new value.

7. Highlight **Save Settings** and press **Enter**.

The new Portal is added to the list.

MAKING iSCSI PORTAL SETTINGS (CLU)

To make portal settings:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Portals** and press **Enter**.

The list of portals displays.

3. Highlight the portal you want to change and press **Enter**.

The portal information screen displays.

4. Highlight **iSCSI Portal Settings** and press **Enter**.
5. Make changes as needed.

- **DHCP** – Enabled or disabled
DHCP is currently supported only for IPv4.
- **ProtocolFamily**– IPv4 or IPv6
- **PrimaryIP** – Primary IP address of this portal
- **PrimaryIPMask** – Subnet mask of this portal
- **Gateway** - Gateway IP address of this portal

- **VlanTag** – VLAN tag number (1 to 4094) for portals associated with a VLAN
 - **TrunkId** – Trunk ID number (1 to 8) for portals associated with a trunk.
6. Highlight **Save Settings** and press **Enter**.
 7. Press Y to acknowledge possible interruption of iSCSI services.
 8. Press Y again to confirm the changes.
 9. Highlight **Return to Previous Menu** and press **Enter** to return to the portal list.

DELETING iSCSI PORTALS (CLU)

To delete an iSCSI portal:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight iSCSI Portals and press **Enter**.

The list of iSCSI portals displays.

3. Highlight the portal you want to delete and press the Spacebar to mark it.
4. Highlight **Delete Marked Targets** and press **Enter**.
5. Press Y to confirm deletion.
6. Press Y again to acknowledge possible interruption of iSCSI services.

The portal is removed from the list.

VIEWING A LIST OF iSCSI SESSIONS (CLU)

To view a list of iSCSI sessions:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Sessions** and press **Enter**.

iSCSI session information includes:

- **ID** – ID number of the session
- **CtrlId** - Session is on this controller
- **TargetAlias** – Alias of the target
- **InitiatorAlias** – Part of the IQN
- **Portal ID** – ID number of the portal
- **Status** – Up or down, active or inactive.

DELETING AN iSCSI SESSION (CLU)

To delete an iSCSI session:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Sessions** and press **Enter**.
3. Highlight the session you want delete and press the spacebar to select it.and press **Enter**.
4. Highlight Delete iSCSI Session and press **Enter**.
5. Press Y to confirm.

VIEWING iSCSI SESSION INFORMATION (CLU)

To view a list of iSCSI sessions:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI Sessions** and press **Enter**.
3. Highlight the session you want and press **Enter**.

iSCSI session information includes:

- | | |
|--|---|
| • Session ID – ID number of the session | • Portal ID – ID number of the portal |
| • Status – Active or inactive | • Target Alias |
| • Initiator Name – SCSI qualified name (iqn) | • Target Name – iSCSI qualified name (iqn) |
| • Portal IP – IP address of the portal | • Initiator IP – IP address of the initiator |
| • Device Type – Initiator or target | • Initiator Source Port – ID number |
| • Target Portal Group – ID number | • ISID – Initiator session ID number |
| • TSIH – Target session identifying handle | • Max Outstanding R2T – Number of PDUs ready to transfer |
| • Execution Throttle – Max number of outstanding commands on any one port | • Max Burst Length – In bytes |
| • Max Rcv Data Seg Length – Receive data segment length | • Default Time to Retain – In seconds |
| • First Burst Length – In bytes | • Initial R2T – Enabled or disabled |
| • Default Time to Wait – In seconds | • Data Digest – Enabled or disabled |
| • Immediate Data – Enabled or disabled | • Data PDU in Order – Enabled or disabled |
| • Header Digest – Enabled or disabled | • Data Seq in Order – Enabled or disabled |
| • CHAP Authentication Type – None, Local, Peer | • Device Access Control – Enabled or disabled |

VIEWING iSCSI iSNS INFORMATION (CLU)

Internet Storage Name Service (iSNS) is a protocol used to facilitate the automated discovery, management, and configuration of iSCSI and Fibre Channel devices on a TCP/IP network.

To view iSNS information:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI iSNS Options** and press **Enter**.

The current iSNS options appear. Information includes:

- **iSNS** – Enabled or disabled
- **iSNSIPAddress** – IP address of the iSNS server
- **iSNSPort** – iSNS port number (1 to 65535) 3205 is the default and recommended number

Items marked with an asterisk (*) are adjustable under Making iSCSI Portal Settings (CLU).

MAKING iSCSI iSNS SETTINGS (CLU)

To make iSNS settings:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI iSNS Options** and press **Enter**.

The current iSNS options appear.

3. Highlight **iSNS Settings** and press **Enter**.
4. Highlight **iSNS** and press the Spacebar to toggle between Enable and Disable.
5. If you chose Enable:
 - **Enter** an IP address.
 - **Enter** a Port number. 3205 is the default and recommended number.

To change a value, highlight the item, press Backspace to delete the current value and type a new value.

6. Highlight **Save Settings** and press **Enter**.
7. Press Y to acknowledge possible interruption of iSCSI services.
8. Press Y again to confirm the changes.
9. Highlight **Return to Previous Menu** and press **Enter** to return to the portal list.

VIEWING A LIST OF iSCSI CHAPs (CLU)

Challenge Handshake Authentication Protocol (CHAP) is an authentication mechanism used to authenticate iSCSI sessions between initiators and targets.

To view a list of iSCSI CHAPs:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI CHAPs** and press **Enter**.

A list of the current CHAPs appears. Information includes:

- **ID** – ID number. Numbering starts at 0.
- **Type** – Peer is one-way. Local is bi-directional.
- **Name** – CHAP name.

ADDING iSCSI CHAPs (CLU)

Verify that CHAP authentication is enabled under Making iSCSI Target Settings (CLU)0.

To add an iSCSI CHAP:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI CHAPs** and press **Enter**.
3. Highlight **Create New CHAP Entry** and press **Enter**.
4. Highlight **Name** and type a name for the CHAP.
5. Highlight **Type** and press the spacebar to toggle between Peer and Local.

Peer is one-way. Local is bi-directional.

6. Highlight **Secret** and type a secret of 12 to 16 characters.
7. Highlight **Retype Secret** and type the secret again to verify.
8. Highlight **Save CHAP Record** and press **Enter**.

The new CHAP is added to the list.

MAKING iSCSI CHAP SETTINGS (CLU)

When you change CHAP settings, you must change the secret. You cannot change the type (peer or local).

To make iSCSI CHAP settings:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI CHAPs** and press **Enter**.
3. Highlight the CHAP you want to edit and press **Enter**.
4. Make changes as needed.
 - Highlight Name and press the backspace key to erase the current value, then type the new value.
 - Highlight **New Secret** and type a secret of 12 to 16 characters.
 - Highlight **Retype New Secret** and type the secret again to verify.
5. Highlight **Save CHAP Record** and press **Enter**.

The edited CHAP appears in the list.

DELETING iSCSI CHAPs (CLU)

To delete an iSCSI CHAP:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **iSCSI CHAPs** and press **Enter**.
3. Highlight the CHAP you want to delete and press **Enter** to mark it.
4. Highlight **Delete Marked Entries** and press **Enter**.
5. Press Y to confirm the deletion.

PINGING A HOST OR SERVER ON THE iSCSI NETWORK (CLU)

This function enables you to ping other network nodes through any one of the Vess R2600’s iSCSI ports.

To ping a host or server on the network:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight Ping and press **Enter**.
3. **Enter** information as required:
 - Highlight **IP address** and type the IP address you want to ping.
 - Highlight **Packet Count** and enter the number of packets you want to send.

- Highlight **Ping Through Controller ID** and choose a controller (1 or 2)
- Highlight **Ping Through Port ID** and choose a port number (1 to 4)

To change a value, highlight the item, press Backspace to delete the current value and type a new value.

4. Highlight Ping and press **Enter**.

The results of the ping are displayed on the iSCSI Ping screen.

VIEWING A LIST OF iSCSI TRUNKS (CLU)

A trunk is the aggregation of two or more iSCSI ports to increase bandwidth.

To view a list of iSCSI trunks:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **Trunk** and press **Enter**.

The list of iSCSI Trunks displays.

- **ID** – ID number of the trunk. Starts at 1.
- **CtrlId** – RAID controller ID, 1 or 2
- **Master Port** – One of the four physical ports on the RAID controller
- **Slave Ports** – Any or all of the remaining physical ports on the same RAID controller
- **Failed Ports** – A slave port that has no iSCSI data connection.
- **State** – Optimal, Sub-Optimal or Failed. Identify and correct the failed iSCSI ports.

ADDING iSCSI TRUNKS (CLU)

Ports must be **enabled** to add them to a trunk. See "Making iSCSI Port Settings (CLU)" on page 148. Vess R2600 supports a maximum of eight trunks.

You cannot use an iSCSI port that has portals configured to it. See "Viewing a List of iSCSI Portals (CLU)" on page 148 and "Deleting iSCSI Portals (CLU)" on page 150.

To add an iSCSI Trunk:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **Trunk** and press **Enter**.
3. Highlight **Create New Trunk** and press **Enter**.
4. **Enter** information as required:
 - Highlight **Controller** and type the controller you want (1 or 2).

- Highlight **Master Port** and type the port number you want (1 to 4).
- Highlight **Slave Ports** and type the port number you want.
For multiple ports, separate the numbers with a comma.
You can choose any or all port numbers except the Master Port number.

5. Highlight **Save Trunk** and press **Enter**.

The new trunk appears in the list.

You can add up to 8 trunks. After you add a trunk, you can assign it to a portal. See "Adding iSCSI Portals (CLU)" on page 149.

MAKING iSCSI TRUNK SETTINGS (CLU)

To make trunk settings:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **Trunk** and press **Enter**.
3. Highlight the trunk you want and press **Enter**.
4. **Enter** information as required:
 - Highlight **Controller** and type the controller you want (1 or 2).
 - Highlight **Master Port** and type the port number you want (1 to 4).
 - Highlight **Slave Ports** and type the port number you want.
For multiple ports, separate the numbers with a comma.
You can choose any or all port numbers except the Master Port number.
5. Highlight **Save Settings** and press **Enter**.

DELETING iSCSI TRUNKS (CLU)

Before you can delete a trunk, you must delete any portals configured on it. See "Deleting iSCSI Portals (CLU)" on page 150.

To delete an iSCSI trunk:

1. From the Main Menu, highlight **iSCSI Management** and press **Enter**.
2. Highlight **Trunk** and press **Enter**.
3. Highlight the trunk you want to delete and press the Spacebar to mark it.
4. Highlight **Delete Marked Trunks** and press **Enter**.
5. Press Y to confirm.

MANAGING BACKGROUND ACTIVITY (CLU)

Background activity refers to any of several functions that take place in the background while normal operation of the Vess R2600 continues.

Background activities work in conjunction with disk arrays and logical drives. See Managing Disk Arrays (CLU) and Managing Logical Drives (CLU) for more information about how and when to use background activities.

Background Activity Management includes the following functions:

- "Viewing Current Background Activities (CLU)" on page 154
- "Making Background Activity Settings (CLU)" on page 154

VIEWING CURRENT BACKGROUND ACTIVITIES (CLU)

From the Main Menu, highlight **Background Activities** and press **Enter**. A count of current background activities appears, including:

- **Rebuild**
- **PDM (Predictive Data Migration)**
- **Synchronization**
- **Redundancy Check**
- **Migration**
- **Transition**
- **Initialization**
- **Media Patrol**

MAKING BACKGROUND ACTIVITY SETTINGS (CLU)

1. From the Main Menu, highlight **Background Activities** and press **Enter**.
2. Highlight **Background Activity Settings** and press **Enter**.
3. Highlight following and press the spacebar to toggle between **Enabled** and **Disabled**.
 - **Media Patrol** – Checks the magnetic media on physical drives
 - **Auto Rebuild** – When enabled and no spare drive is available, the disk array begins to rebuild as soon as you replace the failed physical drive with an unconfigured physical drive of equal or greater size
4. Highlight following and press the spacebar to toggle through **Low**, **Medium**, and **High** rates:
 - **Rebuild** – Rebuilds data to a replacement physical drive in a disk array
 - **Migration** – Change RAID level or add physical drives to disk arrays
 - **PDM** – Migrates data from a suspect physical drive to a replacement drive in a disk array
 - **Transition** – Returns a revertible spare drive to spare status
 - **Synchronization** – Checks the data integrity on disk arrays
 - **Initialization** – Full initialization sets all data bits in the logical drive to a specified pattern, such as all zeros
 - **Redundancy Check** – Checks, reports and can correct data inconsistencies in logical drives

The rates are defined as follows:

- **Low** – Fewer resources to activity, more to data read/write.
- **Medium** – Balance of resources to activity and data read/write.
- **High** – More resources to activity, fewer to data read/write.

5. Highlight the following PDM trigger settings and press the backspace key to erase the current value:
 - **BBM Threshold** – 1 to 2048 reassigned blocks
 - **Media Patrol Threshold** – 1 to 2048 error blocks
6. Press **Control-A** to save your settings.

WORKING WITH THE EVENT VIEWER (CLU)

Working with the Event Viewer includes the following functions:

- "Viewing Runtime Events (CLU)" on page 155
- "Clearing Runtime Events (CLU)" on page 155
- "Viewing NVRAM Events (CLU)" on page 155
- "Clearing NVRAM Events (CLU)" on page 155

The Event Viewer displays log of subsystem events. Events are classified as:

- **Runtime Events** – A list of and information about the 1023 most recent runtime events recorded since the subsystem was started
- **NVRAM Events** – A list of and information about the most important events over multiple subsystem startups. NVRAM events are stored in non-volatile memory

Event Severity Levels	
Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

VIEWING RUNTIME EVENTS (CLU)

To display Runtime Events:

1. From the Main Menu, highlight **Event Viewer** and press **Enter**.

The log of Runtime Events appears. Events are added to the top of the list. Each item includes:

- **Sequence number** – Begins with 0 at system startup.
- **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
- **Severity** – See the table above.
- **Timestamp** – Date and time the event happened.
- **Description** – A description of the event in plain language.

2. Press the up and down arrow keys to scroll through the log.

CLEARING RUNTIME EVENTS (CLU)

To clear the Runtime Event log:

1. From the Main Menu, highlight **Event Viewer** and press **Enter**.
2. Highlight **Clear Runtime Event Log** and press **Enter**.
3. Press Y to confirm.

VIEWING NVRAM EVENTS (CLU)

This screen displays a list of and information about the most important events over multiple subsystem startups.

To display NVRAM events:

1. From the Main Menu, highlight **Event Viewer** and press **Enter**.
2. Highlight **NVRAM Events** and press **Enter**.

The log of NVRAM Events appears. Events are added to the top of the list. Each item includes:

- **Sequence number** – Begins with 0 at system startup.
- **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
- **Severity** – See the table on the previous page.
- **Timestamp** – Date and time the event happened.
- **Description** – A description of the event in plain language.

3. Press the up and down arrow keys to scroll through the log.

CLEARING NVRAM EVENTS (CLU)

To clear the Runtime Event log:

1. From the Main Menu, highlight **Event Viewer** and press **Enter**.
2. Highlight **NVRAM Events** and press **Enter**.
3. Highlight **Clear NVRAM Event Log** and press **Enter**.
4. Press Y to confirm.

WORKING WITH LUN MAPPING (CLU)

LUN Mapping includes the following functions:

- "Enabling LUN Mapping (CLU)" on page 156
- "Viewing a List of Initiators (CLU)" on page 156
- "Adding an Initiator (CLU)" on page 156
- "Viewing a List of LUN Maps (CLU)" on page 157
- "Adding a LUN Map (CLU)" on page 157
- "Editing a LUN Map (CLU)" on page 158
- "Deleting a LUN Map (CLU)" on page 158
- "Changing the Active LUN Mapping Type (CLU)" on page 158

ENABLING LUN MAPPING (CLU)

LUN Mapping must be enabled in order for Vess R2600 to recognize an initiator.

To enable LUN mapping:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Highlight one of the following options and press **Enter**.
 - **LUN Mapping: Intiators**
4. Highlight **Enable LUN Mapping** (Currently DISABLED) and press **Enter**.

A “Logical drives may become invisible” message appears.

5. Press any key to continue.
6. Press Y to confirm.

LUN mapping is enabled.

VIEWING A LIST OF INITIATORS (CLU)

LUN Mapping must be enabled in order for Vess R2600 to recognize an initiator.

To view a list of FC or iSCSI initiators:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Highlight **LUN Mapping: Initiators** and press **Enter**.

A list of the current initiators appears.

ADDING AN INITIATOR (CLU)

You must add an initiator to the Vess R2600’s initiator list in order to use the initiator to create a LUN.

To add an initiator to the Vess R2600’s list:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Highlight **LUN Mapping: Initiators** and press **Enter**.
4. Highlight **Create New Initiator** and press **Enter**.
5. Type a name for the initiator in the field provided.
 - **Fibre Channel** – A Fibre Channel initiator name is the World Wide Port Name of the device and is composed of a series of eight, two-digit hexadecimal numbers.
Example: **10-00-00-00-c9-73-2e-8b**
 - **iSCSI** – An iSCSI initiator name is the iSCSI name of the initiator device and is composed of a single text string.
Example: **iqn.1991-05.com.microsoft:promise-29353b7**
Obtain the initiator name from the initiator utility on your host system.
Note that the initiator name you input must match exactly in order for the connection to work.
6. Highlight **Save Initiator** and press enter.

The new initiator appears in the list.

DELETING AN INITIATOR (CLU)



Caution

If you delete an initiator, you delete the LUN map associated with that initiator. Verify that the LUN map is no longer needed before deleting the initiator

To delete an initiator:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Highlight the initiator you want to delete and press the spacebar to mark it.

The mark is an asterisk (*) to the left of the listing.

4. Highlight **Delete Marked Initiators** and press **Enter**.
5. Press Y to confirm the deletion.

VIEWING A LIST OF LUN MAPS (CLU)

To view a list of LUN maps:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Do one of the following actions:
 - Highlight **LUN Mapping: Initiators** and press **Enter**. Then highlight an initiator and press **Enter**.

The list of logical drives with corresponding LUN maps appears.

ADDING A LUN MAP (CLU)

For FC & iSCSI systems, you can set up an Initiator LUN map.

You can set up the LUN map type on the same subsystem but only one LUN map type can be active at a time.

A maximum of 256 logical drives can be mapped to an FC initiator or to an iSCSI initiator.

To assign a LUN to an FC or iSCSI initiator, add the initiator first. See Adding an Initiator (CLU).

LUN mapping must be enabled in order to map a LUN. See Enabling LUN Mapping (CLU).

MAPPING A LUN TO AN FC INITIATOR

To map a LUN to an FC initiator:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Do the following actions:
 - Highlight **LUN Mapping: Initiators** and press **Enter**. Then highlight an initiator and press **Enter**.
A list of logical drives displays.
4. In the LUN field, press the backspace key to erase the current value, then type the LUN you want to assign to this initiator, from 0 to 255.

Each logical drive can have only one LUN and must have a unique LUN.

If you make a error, press **Control-AR** to restore the current LUN.

5. Press **Control-A** to save the LUN map.

MAPPING A LUN TO AN iSCSI INITIATOR OR TARGET

To map a LUN to an iSCSI initiator or target:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Do one of the following actions:
 - Highlight **LUN Mapping: Initiators** and press **Enter**. Then highlight an initiator and press **Enter**.
A list of logical drives displays.
4. In the LUN field, press the backspace key to erase the current value, then type the LUN you want to assign to this target, from 0 to 255.
Each logical drive can have only one LUN and must have a unique LUN.
If you make a error, press **Control-AR** to restore the current LUN.
5. Press **Control-A** to save the LUN map.

EDITING A LUN MAP (CLU)

Editing a LUN map is the action of assigning a logical drive or LUN to an initiator. By changing the assignment, you change the initiator’s access.

To edit a LUN map:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Do one of the following actions:
 - Highlight **LUN Mapping: Initiators** and press **Enter**. Then highlight an initiator and press **Enter**.
A list of logical drives displays.
4. In the LUN field, press the backspace key to erase the current value, then type the LUN you want to assign to this initiator, from 0 to 255.
Each logical drive can have only one LUN and must have a unique LUN.
If you make a error, press **Control-AR** to restore the current LUN.
5. Press **Control-A** to save the LUN map.

DELETING A LUN MAP (CLU)

Deleting a LUN map prevents the initiator from accessing the LUN while LUN masking is enabled.

To delete a LUN map:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Do one of the following actions:
 - Highlight **LUN Mapping: Initiators** and press **Enter**. Then highlight an initiator and press **Enter**.
A list of logical drives displays.
4. In the LUN field, press the backspace key to erase the current value.
Leave the field blank.
5. Press **Control-A** to save the initiator, port, or target without a LUN map.

CHANGING THE ACTIVE LUN MAPPING TYPE (CLU)

For FC systems, you can set up an Initiator type LUN map.

For iSCSI systems, you can set up an Initiator type LUN map.

You can set up both LUN map types on the same subsystem but only one LUN map type can be active at a time.

To change the active LUN map type:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **LUN Mapping** and press **Enter**.
3. Highlight Active LUN Mapping Type and press the Spacebar to toggle between choices:
 - FC subsystems, choose the Initiator option.
 - iSCSI subsystems, choose the Initiator option.
4. Press **Control-A** to save your setting.

MANAGING UPS UNITS (CLU)

Uninterruptible Power Supply (UPS) Management includes the following functions:

- "Viewing a List of UPS Units (CLU)" on page 159
- "Making UPS Settings (CLU)" on page 159
- "Viewing UPS Information (CLU)" on page 160

VIEWING A LIST OF UPS UNITS (CLU)

To view a list of UPS units supporting the Vess R2600:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **UPS Management** and press **Enter**.

Information in the UPS List includes:

- **Operational Status** – OK means Normal.
On AC means the UPS is connected to a viable external AC power source.
On Battery means the external AC power source is offline and the UPS is running on battery power.
- **Capacity** – Backup capacity expressed as a percentage.
- **Remaining Minutes** – Number of minutes the UPS is expected to power your system in the event of a power failure.
- **Loading** – Actual output of UPS as a percentage of the rated output. See the Note below.



Note

The maximum recommended Loading Ratio varies among models of UPS units. The general range is 60% to 80%. If the reported Loading Ratio exceeds the recommended value for your UPS unit:

Have fewer subsystems or peripherals connected to this UPS unit.
Add more UPS units, or use a higher-capacity UPS unit, to protect your RAID systems.

MAKING UPS SETTINGS (CLU)

These settings control how the Vess R2600 subsystem detects the UPS unit and responds to data reported by the UPS unit.

To make UPS settings:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **UPS Management** and press **Enter**.
3. Highlight **UPS Settings** and press **Enter**.
4. Perform the following actions as required:
 - Verify the Current UPS Communication method. See Note 1:
SNMP – Network connection.
Serial – Serial connection.
Unknown – No connection.
 - Choose a Detection Setting from the drop-down menu:
Automatic – Default. If a UPS is detected when the subsystem boots, the settings changes to Enable.
Enable – Monitors UPS. Settings changes, reports warnings, and logs events.
Disable – Monitors UPS only.
 - Type values into the Threshold fields. See Note 2:
Running Time Remaining Threshold – Actual time below this value resets adaptive writeback cache to writethrough.
Warning Temperature Threshold – Actual temperature above this value triggers a warning and logs an event.
Loading Ratio Threshold – Actual loading ratio (percentage) above this threshold triggers a warning and logs an event. See Note 3.
Battery Charge Remaining Threshold – Reserve capacity below this percentage triggers a warning and logs an event.
 - For UPS units with network cards, type the IP addresses or DNS names in fields UPS 1 and UPS 2. See Note 4.

5. Press **Control-A** to save your settings.

Note 1: Vess R2600 supports multiple UPS units using network or serial connections, but not a combination of both methods.

Note 2: Detection Setting must be set to Auto. If a UPS is detected, the settings changes to Enable.

Note 3: The maximum recommended Loading Ratio varies among models of UPS units. The general range is 60% to 80%.

Note 4: To specify UPS units by DNS names, ask your IT administrator to add the DNS names to the DNS server, before you make UPS settings.

VIEWING UPS INFORMATION (CLU)

To view information about a specific UPS unit:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **UPS Management** and press **Enter**.
3. Highlight the UPS unit you want and press **Enter**.

UPS information includes:

- **UPS ID**
- **Model Name**
- **Serial Number**
- **Firmware Version**
- **Manufacture Date**
- **Voltage Rating** – Output voltage of the UPS.
- **Battery Capacity** – Backup capacity expressed as a percentage.
- **Remaining Backup Time** – Number of minutes the UPS is expected to power your system in the event of a power failure.
- **Loading Ratio** – Actual output of UPS as a percentage of the rated output. See the Note below.
- **Temperature** – Reported temperature of the UPS unit.



Note

The maximum recommended Loading Ratio varies among models of UPS units. The general range is 60% to 80%. If the reported Loading Ratio exceeds the recommended value for your UPS unit:

Have fewer subsystems or peripherals connected to this UPS unit.

Add more UPS units, or use a higher-capacity UPS unit, to protect your RAID systems.

MANAGING USERS (CLU)

User Management includes the following functions:

- "Viewing User Information (CLU)" on page 160
- "Creating a User (CLU)" on page 160
- "Changing Another User's Settings (CLU)" on page 161
- "Changing Your Own User Settings (CLU)" on page 161
- "Changing Another User's Password (CLU)" on page 161
- "Changing Your Own Password (CLU)" on page 162
- "Deleting a User (CLU)" on page 162

VIEWING USER INFORMATION (CLU)

Each user types their user name and password to log into the CLI.

To view a list of current user accounts:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **User Management** and press **Enter**.

A list of the current users appears.

CREATING A USER (CLU)

To create a new user account:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **User Management** and press **Enter**.
3. Highlight **Create New User** and press **Enter**.
4. Highlight each field and type in the appropriate information:
 - **User name** (Maximum 31 characters. Use letters, numbers, and underscore. No spaces.)
 - **Password** (Optional. Maximum 31 characters. Use letters, numbers, and underscore.)
 - **Display name** (Optional)
 - **User's email address**

5. Highlight **Privilege** and press the space bar to toggle though the options.

See the Table on the next page.

6. Press **Control-A** to save the user.

User Privileges	
Level	Meaning
View	Allows the user to see all status and settings but not to make any changes
Maintenance	Allows the user to perform maintenance tasks including Rebuilding, PDM, Media Patrol, and Redundancy Check
Power	Allows the user to create (but not delete) disk arrays and logical drives, change RAID levels, change stripe size; change settings of components such as disk arrays, logical drives, physical drives, and the controller
Super	Allows the user full access to all functions including create and delete users and changing the settings of other users, and delete disk arrays and logical drives. The default “administrator” account is a Super User

CHANGING ANOTHER USER’S SETTINGS (CLU)

The Administrator or a **Super User** can change other users’ settings.

To change user settings:

- From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
- Highlight **User Management** and press **Enter**.
- Highlight the User whose settings you want to change and press **Enter**.
- Highlight **Privilege** and press the space bar to toggle though the options.

See the Table above.
- Highlight **Status** and press the space bar to toggle between **Enabled** and **Disabled**.
- Highlight the items you want and press the backspace key to erase the current value, then type the new value:
 - User name**
 - Email address**
- Press **Control-A** to save the settings.



Important

If a user is logged-in when his account is disabled, the user is immediately logged-out.

CHANGING YOUR OWN USER SETTINGS (CLU)

Each user can change their display name and email address.

To change your user settings:

- From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
- Highlight **User Management** and press **Enter**.
- Highlight your name and press **Enter**.
- Highlight the items you want and press the backspace key to erase the current value, then type the new value:
 - User name**
 - Email address**
- Press **Control-A** to save the settings.

CHANGING ANOTHER USER’S PASSWORD (CLU)

The Administrator or a **Super User** can change other users’ passwords.

To change a password:

- From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
- Highlight **User Management** and press **Enter**.
- Highlight the User whose password you want to change and press **Enter**.
- Highlight **Change Password...** and press **Enter**.
- Highlight **New Password** and type a new password.

Maximum 31 characters. Use letters, numbers, and underscore.
- Highlight **Retype Password** and type the new password again to verify.
- Press **Control-A** to save the new password.



Note

To reset the Administrator’s password to the factory default, see "Restoring Factory Defaults (CLU)" on page 167.

CHANGING YOUR OWN PASSWORD (CLU)

Each user can change their own password.

To change your password:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **User Management** and press **Enter**.
3. Highlight your name and press **Enter**.
4. Highlight **Change Password...** and press **Enter**.
5. Highlight **Old Password** and type your current password.
6. Highlight **New Password** and type a new password.

Maximum 31 characters. Use letters, numbers, and underscore.

7. Highlight **Retype Password** and type the new password again to verify.
8. Press **Control-A** to save the new password.

DELETING A USER (CLU)

The Administrator or a **Super User** can delete other users. You cannot delete the account you used to log in.

There must always be one **Super User** account.

Rather than deleting a user, consider disabling a user account. See Changing Another User’s Settings (CLU) a.

To delete a user:

1. Log in under a user name other than the one you want to delete.
2. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
3. Highlight **User Management** and press **Enter**.
4. Highlight the user you want to delete and press the spacebar to mark it.

The mark is an asterisk (*) to the left of the listing.

5. Highlight **Delete Marked Users** and press **Enter**.
6. Press Y to confirm the deletion.

WORKING WITH SOFTWARE MANAGEMENT (CLU)

Software Management includes the following functions:

- "Making Email Settings (CLU)" on page 163
- "Making SLP Settings (CLU)" on page 163
- "Making Web Server Settings (CLU)" on page 163
- "Making SSH Settings (CLU)" on page 164
- "Making SNMP Settings (CLU)" on page 164
- "Managing SNMP Trap Sinks (CLU)" on page 164
- "Making Netsend Settings (CLU)" on page 165
- "Managing Netsend Recipients (CLU)" on page 165

MAKING EMAIL SETTINGS (CLU)

By default, Email service is set to Automatic and its normal status is Started.

To make Email service settings:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Software Management** and press **Enter**.
3. Highlight **Email** and press **Enter**.
4. Highlight **Startup Type** and press the spacebar to toggle between **Automatic** and **Manual**.
5. Highlight the following and press the backspace key to erase the current value, then type the new value:
 - SMTP server IP address or server name
 - Server Port number (25 is the default)

6. Highlight **Authentication** and press the spacebar to toggle between **Yes** and **No**.

If you selected Yes, type in a User name and Password in the fields provided.

7. The following items are optional but recommended. Highlight and press the backspace key to erase the current value, then type the new value:
 - Sender's email address
 - Subject Line for the email message

8. Press **Control-A** to save your settings.

To start, stop or restart the Email service, highlight **Start**, **Stop** or **Restart** and press **Enter**.

MAKING SLP SETTINGS (CLU)

By default, SLP service is set to Automatic and its normal status is Started.

To make SLP service settings:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Software Management** and press **Enter**.
3. Highlight **SLP** and press **Enter**.
4. Highlight **Startup Type** and press the spacebar to toggle between **Automatic** and **Manual**.
5. Press **Control-A** to save your settings.

To start, stop or restart the SLP service, highlight **Start**, **Stop**, or **Restart** and press **Enter**.

MAKING WEB SERVER SETTINGS (CLU)

By default, Web Server service is set to Automatic and its normal status is Started.

To make Web Server service settings:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Software Management** and press **Enter**.
3. Highlight and press **Enter**.
4. Highlight **Startup Type** and press the spacebar to toggle between **Automatic** and **Manual**.
5. Highlight the following and press the backspace key to erase the current value, then type the new value:
 - HTTP Port (80 is the default)
 - Session Time Out (24 minutes is the default. 1440 minutes = 24 hours)
6. Highlight **SSL** and press the spacebar to toggle between **Enabled** and **Disabled**.
7. Highlight **HTTPS Port** and press the backspace key to erase the current value, then type the new value. 443 is the default.
8. Press **Control-A** to save your settings.

To start, stop or restart the service, highlight **Start**, **Stop**, or **Restart** and press **Enter**.

MAKING TELNET SETTINGS (CLU)

By default, Telnet service is set to Manual and its normal status is Stopped.

To make Telnet service settings:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Software Management** and press **Enter**.
3. Highlight **Telnet** and press **Enter**.
4. Highlight **Startup Type** and press the spacebar to toggle between **Automatic** and **Manual**.
5. Highlight the following and press the backspace key to erase the current value, then type the new value:
 - Port number (2300 is the default)
 - Session Time Out (24 minutes is the default. 1440 minutes = 24 hours)
 - Maximum number of connections (4 is the default)
6. Press **Control-A** to save your settings.

To start, stop or restart the Telnet service, highlight **Start**, **Stop**, or **Restart** and press **Enter**.

MAKING SSH SETTINGS (CLU)

By default, Secure Shell (SSH) service is set to Automatic and its normal status is Started.

To make SSH settings:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Software Management** and press **Enter**.
3. Highlight **SSH** and press **Enter**.
4. Highlight **Startup Type** and press the spacebar to toggle between **Automatic** and **Manual**.
5. Highlight the following and press the backspace key to erase the current value, then type the new value:
 - Port number (22 is the default)
 - Session Time Out (24 minutes is the default. 1440 minutes = 24 hours)
 - Maximum number of connections (4 is the default)
6. Press **Control-A** to save your settings.

MAKING SNMP SETTINGS (CLU)

By default, Simple Network Management Protocol (SNMP) service is set to Automatic and its normal status is Started.

To make SNMP service settings:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Software Management** and press **Enter**.
3. Highlight **SNMP** and press **Enter**.
4. Highlight **Startup Type** and press the spacebar to toggle between **Automatic** and **Manual**.
5. Highlight the following and press the backspace key to erase the current value, then type the new value:
 - **Port Number** – 161 is the default
 - **System Name** – (optional) Type a system name in this field
 - **System Location** – Type a country name in this field
 - **System Contact** – Type the email address of your system administrator in this field
 - **Read Community** – Type a community name in this field
 - **Write Community** – private (no change possible)
6. Press **Control-A** to save your settings.

To start, stop or restart the SNMP service, highlight **Start**, **Stop**, or **Restart** and press **Enter**.

MANAGING SNMP TRAP SINKS (CLU)

VIEWING A LIST OF TRAP SINKS

To create a trap sink:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Software Management** and press **Enter**.
3. Highlight **SNMP** and press **Enter**.
4. Highlight **Trap Sinks** and press **Enter**.

A list of the current trap sinks appears.

ADDING A TRAP SINK

- To add a trap sink:
1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
 2. Highlight **Software Management** and press **Enter**.
 3. Highlight **SNMP** and press **Enter**.
 4. Highlight **Trap Sinks** and press **Enter**.
 5. Highlight **Create New Trap Sink** and press **Enter**
 6. Highlight **Trap Sink IP address** and press the backspace key to erase the current value, then type the new IP address in this field.
 7. Highlight **Trap Filter** and press the spacebar to toggle through the severity levels.

- See the Table below.
8. Press **Control-A** to save the Trap Sink.

Event Severity Levels	
Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

DELETING A TRAP SINK

- To delete a trap sink:
1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
 2. Highlight **Software Management** and press **Enter**.
 3. Highlight **SNMP** and press **Enter**.
 4. Highlight **Trap Sinks** and press **Enter**.
 5. Highlight the trap sink you want to delete and press the spacebar to mark it.

- The mark is an asterisk (*) to the left of the listing.
6. Highlight **Delete Marked Entries** and press **Enter**.

MAKING NETSEND SETTINGS (CLU)

- By default, Netsend service is set to Manual and its normal status is Stopped.
- To make Netsend service settings:
1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
 2. Highlight **Software Management** and press **Enter**.
 3. Highlight **Netsend** and press **Enter**.
 4. Highlight **Startup Type** and press the spacebar to toggle between **Automatic** and **Manual**.
 5. Press **Control-A** to save your settings.
- To start, stop or restart the Netsend service, highlight **Start**, **Stop**, or **Restart** and press **Enter**.

MANAGING NETSEND RECIPIENTS (CLU)

Vess R2600’s Netsend service sends Vess R2600 subsystem events in the form of text messages to your Host PC and other networked PCs.

NETSEND REQUIREMENTS

- In order to use Netsend:
- NetSend must be running the Vess R2600
 - You must provide the IP address for each recipient PC
 - The Messenger service must be running on each recipient PC
- If your Netsend and Messenger service settings are correct but the recipient PC does not receive event messages, check the recipient PC’s Firewall settings. Refer to your OS documentation for more information.

ADDING NETSEND RECIPIENTS

- To add a Netsend recipient:
1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
 2. Highlight **Software Management** and press **Enter**.
 3. Highlight **Netsend** and press **Enter**.
 4. Highlight **Message Recipients** and press **Enter**.
 5. Highlight **Create New Message Recipient** and press **Enter**.

- 6. Type the recipient’s IP address into the field provided.
- 7. Highlight **Message Event Severity Filter** and press the spacebar to change severity levels.

The selected level and all higher severity levels of severity are reported.

See the Table below.

- 8. Press **Control-A** to save your settings.

Event Severity Levels	
Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

DELETING NETSEND RECIPIENTS

To delete a Netsend recipient:

- 1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
- 2. Highlight **Software Management** and press **Enter**.
- 3. Highlight **Netsend** and press **Enter**.
- 4. Highlight **Message Recipients** and press **Enter**.
- 5. Highlight the recipient you want to delete and press the spacebar to mark it.

The mark is an asterisk (*) to the left of the listing

- 6. Highlight **Delete Marked Entries** and press **Enter**.

FLASHING THROUGH TFTP

Use this function to flash (update) the firmware on the Vess R2600. See "Updating with the CLU" on page 173

VIEWING FLASH IMAGE INFORMATION (CLU)

Flash image information refers to the package of firmware components running on your Vess R2600 controller or controllers.

To view flash image information:

- 1. From the Main Menu, highlight **Additional Info and Management**, and press **Enter**.
- 2. Highlight **Flash Image Version Info** and press **Enter**.

The flash image information displays on the screen:

- Enclosure Number – 1 (one) is the Head Unit. Other numbers are cascaded or expanded subsystems
- Running Image Info – Firmware currently running on the controllers
- FlashedImage Info – Firmware flashed to memory
- Image Type – A specific component
- Controller ID – 1 or 2
- Version number
- Build date
- Flash (installation) date

If the Running and FlashedImage do not match, the Vess R2600 has not restarted since the firmware was last updated. Restart the Vess R2600 to run the FlashedImage firmware package. See "Restarting a Subsystem" on page 176

Note that all of these components are upgraded together in a package. See "Updating with the CLU" on page 173.

CLEARING STATISTICS (CLU)

This function clears the statistical counts for the RAID controller, Fibre Channel ports, iSCSI ports, physical drives, and logical drives. To clear statistics:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Clear Statistics** and press **Enter**.
3. Press Y to confirm the deletion.

RESTORING FACTORY DEFAULTS (CLU)

This function restores the factory default settings to the firmware and software items you select.



Caution

Restoring default settings can disrupt your Vess R2600 functions. Use this feature only when necessary.

If you restore Management Network settings, you lose your network connection to the Vess R2600.



Note

To reset the Administrator's password to the factory default, see "Restoring Factory Defaults (CLU)" on page 167.

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Restore Factory Defaults** and press **Enter**.
3. Highlight the setting groups you want to restore and press the spacebar to toggle between **Yes** and **No**.

Yes means this setting is restored to the default value.

No means the current setting remains untouched.

4. Highlight **Restore Factory Defaults** and press **Enter**.
5. Press Y to confirm the reset.

SHUTTING DOWN THE SUBSYSTEM (CLU)

There are two methods for shutting down the subsystem. Choose one of the following procedures:

- "Shutting down the Vess R2600 – Telnet Connection" on page 168
- "Shutting down the Vess R2600 – Serial Connection" on page 168

SHUTTING DOWN THE VESS R2600 – TELNET CONNECTION

This function shuts down the Vess R2600 subsystem on a Telnet connection. Additional action is required, as described below.



Important

If you have a JBOD Expansion, always power off the RAID subsystem first. Then power off the JBOD subsystems.

To shutdown the RAID subsystem:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Shutdown or Restart** and press **Enter**.
3. Highlight **Option** and press the spacebar to display **Shutdown**.
4. Highlight **Submit** and press **Enter**.

A warning message appears.

5. Press Y to continue.

The screen goes blank.

6. Wait for no less than two minutes.
7. The subsystem will then turn off the system power one after another.

SHUTTING DOWN THE VESS R2600 – SSH CONNECTION

This function shuts down the Vess R2600 subsystem on a SSH connection. Additional action is required, as described below.



Important

If you have a JBOD Expansion, always power off the RAID subsystem first. Then power off the JBOD subsystems.

To shutdown the RAID subsystem:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Shutdown or Restart** and press **Enter**.
3. Highlight **Option** and press the spacebar to display **Shutdown**.
4. Highlight **Submit** and press **Enter**.

A warning message appears.

5. Press Y to continue.
6. Close your SSH session.
7. Wait for no less than two minutes.
8. The subsystem will then turn off the system power one after another.

SHUTTING DOWN THE VESS R2600 – SERIAL CONNECTION

This function shuts down the Vess R2600 subsystem on a serial connection. Additional action is required, as described below.



Important

If you have a JBOD Expansion, always power off the RAID subsystem first. Then power off the JBOD subsystems.

To shutdown the RAID subsystem:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Shutdown or Restart** and press **Enter**.
3. Highlight **Shutdown or Restart** and press **Enter**.
4. Highlight **Option** and press the spacebar to display **Shutdown**.

5. Highlight **Submit** and press **Enter**.

A warning message appears.
6. Press Y to continue.
7. The subsystem will then turn off the system power one after another.

STARTING UP AFTER SHUTDOWN

There are two methods for shutting down the subsystem. Choose one of the following procedures:

- "Starting up the Vess R2600 – Telnet Connection" on page 169
- "Starting up the Vess R2600 – SSH Connection" on page 169
- "Starting up the Vess R2600 – Serial Connection" on page 169

STARTING UP THE VESS R2600 – TELNET CONNECTION

To start the RAID subsystem:

1. Manually turn on the system by pressing the power button on the front left side.
2. Wait about two minutes.
3. Establish a Telnet connection to the Vess R2600.

See Making a Telnet Connection.

If you cannot log in, wait 30 seconds and try again.

4. Type **menu** and press **Enter** to open the CLU.

STARTING UP THE VESS R2600 – SSH CONNECTION

To start the RAID subsystem:

1. Manually turn on the system by pressing the power button on the front left side.
2. Wait about two minutes.
3. Establish a SSH connection to the Vess R2600.

See Making a SSH Connection.

If you cannot log in, wait 30 seconds and try again.

4. Type **menu** and press **Enter** to open the CLU.

STARTING UP THE VESS R2600 – SERIAL CONNECTION



Important

If you have a JBOD Expansion, always power on the JBOD subsystems first. Then power on the RAID subsystem.

To start the RAID subsystem:

1. Manually turn on the system by pressing the power button on the front left side
2. Wait about two minutes
3. Establish a serial connection to the Vess R2600.

See Making a Serial Connection.

When the Login: prompt appears, the start up is finished.

4. Type **menu** and press **Enter** to open the CLU.

RESTARTING THE SUBSYSTEM

There are two methods for restarting the subsystem. Choose one of the following procedures:

- "Restarting Vess R2600 – Telnet Connection" on page 170
- "Restarting Vess R2600 – SSH Connection" on page 171
- "Restarting Vess R2600 – Serial Connection" on page 171



Note

If you have a JBOD Expansion, you are not required to restart the JBOD subsystems when you restart the RAID subsystem.

-

RESTARTING VESS R2600 – TELNET CONNECTION

To restart the RAID subsystem:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Shutdown or Restart** and press **Enter**.
3. Highlight **Option** and press the spacebar to display **Restart**.
4. Highlight **Submit** and press **Enter**.

A warning message appears.

5. Press Y to continue.

The screen goes blank.

6. Wait about two minutes.
7. Re-establish your Telnet connection to the Vess R2600 CLU.

See Making a Telnet Connection.

If you cannot re-establish a connection, wait 30 seconds and try again.

RESTARTING VESS R2600 – SSH CONNECTION

To restart the RAID subsystem:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Shutdown or Restart** and press **Enter**.
3. Highlight **Option** and press the spacebar to display **Restart**.
4. Highlight **Submit** and press **Enter**.

A warning message appears.

5. Press Y to continue.
6. Close your SSH session.
7. Wait about two minutes.
8. Re-establish your SSH connection to the Vess R2600 CLU.

See Making a SSH Connection.

If you cannot re-establish a connection, wait 30 seconds and try again.

RESTARTING VESS R2600 – SERIAL CONNECTION

To restart the RAID subsystem:

1. From the Main Menu, highlight **Additional Info and Management** and press **Enter**.
2. Highlight **Shutdown or Restart** and press **Enter**.
3. Highlight **Option** and press the spacebar to display **Restart**.
4. Highlight **Submit** and press **Enter**.

A warning message appears.

5. Press Y to continue.

The screen displays shutdown and startup functions.

6. When the **Login:** prompt appears, log into the CLU again.

BUZZER

MAKING BUZZER SETTINGS

The buzzer sounds to inform you that the Vess R2600 needs attention. See "Vess R2600 is Beeping" on page 201 for more information.

To make buzzer settings:

1. From the Main Menu, highlight **Buzzer** and press **Enter**.

A list of Controllers appears with the current buzzer setting and status.

2. Highlight the Controller whose buzzer you want to set and press **Enter**.
3. Highlight **Enabled** and press the spacebar to toggle between **Yes** and **No**.
4. Press **Control-A** to save your settings.

SILENCING THE BUZZER



Caution

This action disables the buzzer for all events.

To silence the buzzer, follow the procedure above for disabling the buzzer.

CHAPTER 6: MAINTENANCE

This chapter covers the following topics:

- "Updating the Subsystem Firmware" on page 172
- "Updating Physical Drive Firmware" on page 175
- "Replacing a Power Supply" on page 177
- "Replacing a Cache Backup Battery" on page 179
- "Replacing a RAID Controller – Dual Controllers" on page 180
- "Replacing a RAID Controller – Single Controller" on page 181

UPDATING THE SUBSYSTEM FIRMWARE

This procedure applies to Vess R2600 RAID subsystems and Vess R2600 JBOD expansion units managed by a Vess R2600 RAID subsystem. There are two methods:

- "Updating with WebPAM PROe" on page 172
- "Updating with USB Support" on page 174

UPDATING WITH WEBPAM PROE

Download the latest firmware image file from PROMISE support:
<http://www.promise.com/support/> and save it to your Host PC or TFTP server.



Important

Verify that no background activities are running on the RAID subsystem.

To update the firmware on the RAID subsystem and JBOD expansion units:

1. Click the **Administration** tab.
2. Click the Firmware **Update** icon.
3. Click the **Controller Firmware Update** tab.

The Controller Firmware Update screen appears showing the current Image Version Number and Build Date.

4. Choose a download option:
 - **Local File through HTTP** – Click the **Browse** button, locate the firmware image file, click the file to choose it, then click the **Open** button.
 - **TFTP Server** – **Enter** the TFTP Server host name or IP address, port number and file name.
5. Optional. Check the Non-disruptive Image Update (NDIU) box.

NDIU updates the RAID controllers and I/O modules one at a time, enabling I/O operations continue during the firmware update. Updates with this option take a longer period of time to complete. All Vess R2600 models support this feature.

6. Click the **Next** button.

The next screen shows the Flash Image (firmware image file) Version Number and Build Date.

7. Click the **Submit** button.

8. In the **Confirmation** box, type the word “**confirm**” in the field provided and click the **Confirm** button.

The progress of the update displays.



Warning

Do NOT power off the RAID subsystem during the update!
Do NOT move to any other screen until the firmware update operation is completed!

When the update is completed a message tells you to reboot the subsystem,

9. Click the **OK** button.

- If you chose the Disruptive Flash Method, the RAID subsystem and JBOD expansion units automatically restart.
- If you chose the Non-Disruptive Flash Method, the system automatically flashes and restarts the RAID controllers one at a time.

AUTOMATIC RESTART

If you did NOT check the NDIU box, the RAID subsystem and JBOD expansion units automatically restart. That action temporarily disrupts I/O operations and drops your WebPAM PROe connection.

To reestablish your WebPAM PROe connection:

1. Wait no less than two minutes.
2. Click **Logout** in the WebPAM PROe Header, then log in again.

If you cannot log in, wait 30 seconds and try again.

3. In your browser, click Logout in the WebPAM PROe Header, then log in again.

If you cannot log in immediately, wait 30 seconds and try again.

UPDATING WITH THE CLU

Download the latest firmware image file from PROMISE support:

<http://www.promise.com/support/> and save it to your Host PC or TFTP server.



Important

Verify that no background activities are running on the RAID subsystem.

To update the firmware on the RAID subsystem and JBOD expansion units:

1. From the Main Menu, highlight **Additional Info and Management**, and press **Enter**.
2. Highlight **Flash through TFTP** and press **Enter**.
3. Highlight **TFTP Server** and type the IP address of your TFTP server in the field provided.
4. Highlight **Port Number** and press the backspace key to erase the current value, then type the new value. 69 is the default.

A list of the current users appears.

5. Highlight **File Name** and type the file name of the firmware image file in the field provided.
6. Highlight **Flash Method** and press the spacebar to toggle between:
 - **Disruptive** – Updates the RAID controllers and I/O modules simultaneously. I/O operations stop during the firmware update.
 - **Non Disruptive** – (NDIU) Updates the RAID controllers and I/O modules one at a time, enabling I/O operations continue during the firmware update. Updates with this option take a longer period of time to complete. All Vess R2600 models support this feature.
7. Highlight **Start** and press **Enter**.



Warning

Do NOT power off the RAID subsystem during the update!
Do NOT move to any other screen until the firmware update operation is completed!

- If you chose the Disruptive Flash Method, the RAID subsystem and JBOD expansion units automatically restart.
- If you chose the Non-Disruptive Flash Method, the system automatically flashes and restarts the RAID controllers one at a time.

AUTOMATIC RESTART

If you chose the Disruptive Flash Method, the RAID subsystem and JBOD expansion units automatically restart. That action temporarily disrupts I/O operations and drops your CLU connection.

After the screen goes blank, wait about two minutes, then re-establish your Telnet connection to the CLU. If you cannot re-establish a connection, wait 30 seconds and try again.


UPDATING WITH USB SUPPORT

USB support uses the disruptive flash method only. Both RAID controllers and all JBOD I/O modules are updated at the same time and momentarily go offline when the RAID subsystem and JBOD unit reboot.

This procedure requires a USB flash device:

- Formatted to FAT 32
- At least 50 MB of free space

Download the latest OPAS_xxxx.zip firmware image file from PROMISE support: <http://www.promise.com/support/> and save it the root folder of the USB flash device.




Important

Verify that no background activities are running on the RAID subsystem.

To update the subsystem firmware using Vess R2600’s USB Support feature:

1. Insert the USB flash device into the bottom USB port on the front panel which is for the RAID controller.

Figure 6 - 1: Vess R2600fiD controller LEDs



The controller status LED blinks green in half-second intervals.

2. Wait until the controller activity LED stops blinking green and starts blinking amber.



Warning

Do NOT power off the RAID subsystem during the update!

Do NOT remove your USB flash device until the LED changes color!

3. Within 30 seconds, remove the USB flash device, then insert the USB flash device back into the same RAID controller.

The remove and insert action confirms that you want to update the firmware.

You can insert the USB flash device back into either USB port but it must be the same RAID controller as step 1.

4. Wait until the controller activity LED displays steady green.
5. Remove the USB flash device.

AUTOMATIC RESTART

After you remove the USB flash device from the RAID controller, the RAID subsystem and any JBOD expansion units automatically restart. That action temporarily disrupts I/O operations and drops your WebPAM PROe or CLU connection.

To reestablish your WebPAM PROe connection:

1. Wait no less than two minutes.
2. Click **Logout** in the WebPAM PROe Header, then log in again.

If you cannot log in, wait 30 seconds and try again.

To reestablish your CLU connection:

After the screen goes blank, wait about two minutes, then re-establish your Telnet connection to the CLU. If you cannot re-establish a connection, wait 30 seconds and try again.

If you have a serial connection to the RAID subsystem, the connection remains during the shut-down and restart. No reconnect is required.

FAILED UPDATE

If the firmware update fails, the controller status LED displays red. See Vess R2600fiD controller LEDs, Vess R2600fiD controller LEDs

1. Remove the USB flash device.
2. Insert the USB flash device into a USB port on your PC.
3. Go to the **OPAX_XXXXXX** folder to obtain the report and log.

Possible causes for an update failure include:

- Less than 50 MB free space on the USB flash device.
- The Vess R2600 firmware image is invalid.
- A background activity is running.

See "Contacting Technical Support" on page 228.

UPDATING PHYSICAL DRIVE FIRMWARE

This feature applies only to PROMISE-supported physical drives. For a list of supported drives, go to PROMISE support: <http://www.promise.com/support/>.

If you have physical drives in your RAID system that are not PROMISE-supported, follow the firmware update procedure from the drive manufacturer.

WebPAM PROE

Download the latest firmware image file from PROMISE support:
<http://www.promise.com/support/> and save it to your Host PC or TFTP server.

To update the firmware on PROMISE-supported physical drives:

1. Click the **Administration** tab.
2. Click the **Firmware Update** icon.
3. Click the **PD Firmware Update** tab.
4. Choose a download option:
 - **Local File through HTTP** – Click the **Browse** button, locate the firmware image file, click the file to choose it, then click the **Open** button.
 - **TFTP Server** – **Enter** the TFTP Server host name or IP address, port number and file name.
5. Click the **Next** button.
6. Click the **Submit** button.

The progress of the update displays.



Warning

Do NOT power off the RAID subsystem during the update!
Do NOT move to any other screen until the firmware update operation is completed!

When the update is completed a message tells you to reboot the subsystem.

7. Click the **OK** button.

Restart the RAID subsystem. See “Restarting a Subsystem” on the next page.

RESTARTING A SUBSYSTEM

This function shuts down the subsystem and then restarts it.



Important

Do NOT turn off the power supply switches on the RAID subsystem or JBOD expansion units.

To restart the subsystem:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Shutdown/Restart** button.
4. Click the **Restart** button.
5. Type the word “**confirm**” in the field provided.
6. Click the **Confirm** button.

When the controller shuts down, your WebPAM PROe connection is lost.

7. Wait no less than two minutes.
8. In your browser, click Logout in the WebPAM PROe Header, then log in again.

If you cannot log in immediately, wait 30 seconds and try again.

REPLACING A POWER SUPPLY


The power supply and its fans are replaced as one unit. There are no individually serviceable parts. No tools are required for this procedure.

REMOVING THE OLD POWER SUPPLY

To remove the power supply:

1. Verify that the PSU status LED is amber or red.
2. Unplug the power cord.
3. Turn the set screw counter-clockwise to loosen it. The screw is retained on the power supply housing.
4. Grasp the handle and pull the power supply straight out of the enclosure.

Figure 6 - 2: Power supply for Vess R2600




INSTALLING A NEW POWER SUPPLY

To install the power supply:

1. Carefully slide the power supply into the enclosure.
2. Turn the set screw clockwise to tighten, DO NOT over tighten.
3. Plug in the power cord.
4. Switch on the power supply.
5. Verify that the new power supply LED is green. See Power supply for Vess R26000

Figure 6 - 3: Insert fresh PSU



REPLACING A COOLING UNIT

A failed Cooling Unit on the Vess R2600 or Vess J2600 can be hot swapped if the other cooling unit is functioning properly (indicated by a green Fan Status LED). Follow the instructions below to replace a problem Cooling Unit.



Important

In the event of a Cooling Unit failure, **DO NOT** remove the failed unit until there is a replacement available and on hand. A single functioning Cooling Unit is adequate for cooling the system as long as the failed Cooling Unit remains in place. Removing the failed unit without replacing it will adversely affect airflow within the enclosure resulting in critical overheating and shutdown of the enclosure.

TO REMOVE THE COOLING UNIT


1. Check the Fan Status LED, a red LED indicates the fan has failed and needs to be replaced, an amber LED indicates a problem.
2. Use a No. 1 Phillips screwdriver to turn each set screw counter-clockwise to loosen them. The screws remain attached to the Cooling Unit so they will not be lost.
3. Pull the detached Cooling Unit out of the subsystem enclosure. You can grasp the posts that house the set screws initially to move the unit out.

TO INSERT A NEW COOLING UNIT

1. Align the fresh Cooling Unit to be received in the empty cooling unit bay of the subsystem enclosure.
2. Slide the Cooling Unit into place until the set screws are able to be reattached to the backplane of the enclosure.
3. Turn the set screws to tighten them.
4. Check the Fan Status LED on the Cooling Unit to make sure it is green (functioning properly).

Figure 6 - 5: Cooling Unit LED indicates a problem with the fan

An amber fan status LED indicates there is a problem. This unit might need to be replaced.



REPLACING A CACHE BACKUP BATTERY

The cache backup battery, also called a Battery Backup Unit (BBU) powers the cache to preserve data that has not been written to the physical drives. The battery is attached to the top of the Cooling Unit assembly. Each Cooling unit has one battery.

To replace a backup battery, swap out the Cooling Unit with the BBU that needs replacement. Follow the instructions in "Replacing a Cooling Unit" on page 178.



Important

DO NOT remove a Cooling Unit unless there is a replacement available and on hand. Removing a single Cooling Unit without replacing it right away will adversely affect airflow within the enclosure resulting in critical overheating and shutdown of the enclosure.



Cautions

Try reconditioning the battery before you replace it. See "Reconditioning a Battery" on page 62 or "Reconditioning a Battery (CLU)" on page 128.

The battery assembly is replaced as a unit. Do not attempt to disconnect the battery by itself. Installing the wrong replacement battery can result in an explosion.


Dispose of used batteries according to the instructions that accompany the battery.

While the battery is removed, your system is vulnerable to data loss if the power fails while data is being written to the logical drives.

If power service has failed, do not remove the battery if the RAID controller's dirty cache LED is flashing. See "Controller LEDs" on page 32.

Figure 6 - 6: Number 2 Cooling Unit with Cache Backup Battery (BBU) LED indicating battery failure

Red or amber BBU LED indicates a battery problem



REPLACING A RAID CONTROLLER – DUAL CONTROLLERS

The RAID controller monitors and manages the logical drives. When the RAID controller is replaced, all of your logical drive data and configurations remain intact because logical drive information is stored on the physical drives.



Important

Do not replace the RAID controller based on LED colors alone. Only replace the RAID controller when directed to do so by PROMISE Technical Support.

The firmware version and amount of SDRAM must be the same on the replacement RAID controller and the other RAID controller in the subsystem.

To obtain firmware and SDRAM information for an installed RAID controller, in WebPAM PROe, click the Administration button then click the Image Version icon.

Replacement RAID controllers do not come with a BBU. Remove the BBU from the old controller and install it into the new one. See Replacing a Cache Backup Battery.



Note



On subsystems with dual RAID controllers, you can hot-swap a controller while the subsystem is running.

REMOVING THE OLD CONTROLLER

To remove a RAID controller:

1. Disconnect all attached cables from the RAID controller,
 - Fibre Channel cables
 - iSCSI cables
 - SAS expansion cables
 - Management port cables
 - Serial cable
 - UPS control cable
2. On the controller handle, squeeze the release tab and pull the handle outward.
3. Pull the RAID controller out of the subsystem enclosure.

Figure 6 - 4: RAID controller release handle and pull out



INSTALLING THE NEW CONTROLLER

To install the new RAID controller:

1. Carefully slide the RAID controller into the enclosure.
2. Gently swing the handle in and press the handle until it locks.
3. Reconnect all cables that were attached to the RAID controller.
 - Fibre Channel cables
 - iSCSI cables
 - SAS expansion cables
 - Management port cables
 - Serial cable
 - UPS control cable

If one of the controllers goes into maintenance mode, see "Maintenance Mode" on page 211.

REPLACING A RAID CONTROLLER – SINGLE CONTROLLER

The RAID controller monitors and manages the logical drives. When the RAID controller is replaced, all of your logical drive data and configurations remain intact because logical drive information is stored on the physical drives.



Caution

The RAID controller is NOT hot-swappable if your Vess R2600 has only one controller. Power-down the Vess R2600 before removing it.



Important

Do not replace the RAID controller based on LED colors alone. Only replace the RAID controller when directed to do so by PROMISE Technical Support. See page 435.



Important

The firmware on the replacement RAID controller must be the same version as the original RAID controller or a later version.

The amount of SDRAM in the replacement RAID controller must be the same as the original RAID controller or greater.

To obtain firmware and SDRAM information for the currently installed RAID controller, click the Administration button then click the Image Version icon.

REMOVING THE OLD CONTROLLER

To remove the RAID controller:

1. Shutdown the Vess R2600.
2. Disconnect all attached cables from the RAID controller,
 - Fibre Channel cables
 - iSCSI cables
 - SAS expansion cables
 - Management port cables
 - Serial cable
 - UPS control cable
3. On the controller handle, squeeze the release tab and pull the handle outward. See RAID controller release handle and pull out, RAID controller release handle and pull out
4. Pull the RAID controller out of the subsystem enclosure.

INSTALLING THE NEW CONTROLLER

To install the new RAID controller:

1. Carefully slide the RAID controller into the enclosure.
2. Gently swing the handle in and press the handle until it locks.
3. Reconnect all cables that were attached to the RAID controller.
 - Fibre Channel cables
 - iSCSI cables
 - SAS expansion cables
 - Management port cables
 - Serial cable
 - UPS control cable
4. Press the power power.

The Vess R2600 restarts.

5. Log into the Vess R2600.

CHAPTER 7: TECHNOLOGY BACKGROUND

This chapter covers the following topics:

- "Disk Arrays" on page 183
- "Logical Drives" on page 184
- "Spare Drives" on page 193
- "RAID Controllers" on page 194
- "iSCSI Management" on page 197
- "Internet Protocols" on page 200

DISK ARRAYS

Disk array technology includes:

- "Media Patrol" on page 183

MEDIA PATROL

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. Media Patrol checks all physical drives assigned to disk arrays and spare drives. Media Patrol does not check unconfigured drives.

Media Patrol checks are enabled by default on all disk arrays and spare drives. You can disable Media Patrol in the disk array and spare drive settings, however that action is not recommended.

Unlike Synchronization and Redundancy Check, Media Patrol is concerned with the condition of the media itself, not the data recorded on the media. If Media Patrol encounters a critical error, it triggers PDM, if PDM is enabled on the disk array.

Media Patrol has three status conditions:

- **Running** – Normal. You can access your logical drives at any time.
- **Yield** – Temporary pause while a read/write operation takes place.
- **Paused** – Temporary pause while another background runs. Or a pause initiated by the user.

PDM

Predictive Data Migration (PDM) is the migration of data from the suspect physical drive to a spare drive, similar to rebuilding a logical drive. But unlike Rebuilding, PDM constantly monitors your physical drives and automatically copies your data to a spare drive **before** the physical drive fails and your logical drive goes Critical.

The following actions trigger PDM:

- A physical drive with unhealthy status (see below)
- Media Patrol finds a critical error
- You initiate PDM manually

PDM also counts the number of media errors reported by Media Patrol. A disk drive becomes unhealthy when:

- A SMART error is reported
- The bad sector remapping table fills to the specified level.

Because data would be lost if written to a bad sector, when a bad sector is detected, the physical drive creates a map around it. These maps are saved in the **bad sector remapping table**, which has a capacity of 512 reassigned blocks and 2048 error blocks

You can specify the maximum levels for the reassigned and error blocks in PDM settings. When the table fills to a specified value, PDM triggers a migration of data from the suspect drive (the disk drive with the bad sectors) to a replacement physical drive.

During data migration, you have access to your logical drives but they respond more slowly to read/write tasks because of the additional operation. The time required for data migration depends on the size of the physical drives.

PDM is enabled on all disk arrays by default. You can disable PDM in the disk array settings, however that action is not recommended.

LOGICAL DRIVES

Logical drive technology includes:

- "RAID Levels" on page 184
- "RAID Level Migration" on page 189
- "Stripe Size" on page 192
- "Sector Size" on page 192
- "Initialization" on page 192
- "Partition and Format" on page 192

RAID LEVELS

RAID (Redundant Array of Independent Disks) allows multiple physical drives to be combined together in a disk array. Then all or a portion of the disk array is formed into a logical drive. The operating system sees the logical drive as a single storage device, and treats it as such.

RAID 0 – STRIPE

When a logical drive is striped, the read and write blocks of data are interleaved between the sectors of multiple physical drives. Performance is increased, since the workload is balanced between drives or “members” that form the logical drive. Identical drives are recommended for performance as well as data storage efficiency.

The disk array’s data capacity is equal to the number of disk drive members multiplied by the smallest drive’s capacity. For example, one 100 GB and three 120 GB drives form a 400 GB (4 x 100 GB) disk array instead of 460 GB.

If physical drives of different capacities are used, there is unused capacity on the larger drives.

RAID 0 logical drives on Vess R2600 consist of one or more physical drives.

Advantages	Disadvantages
<ul style="list-style-type: none">• Implements a striped disk array, the data is broken down into blocks and each block is written to a separate disk drive• I/O performance is greatly improved by spreading the I/O load across many channels and drives• No parity calculation overhead is involved	<ul style="list-style-type: none">• Not a true RAID because it is not fault-tolerant• The failure of just one drive results in all data in an disk array being lost• Should not be used in mission critical environments

Recommended Applications for RAID 0:

- Image Editing
- Pre-Press Applications
- Any application requiring high bandwidth

RAID 1 – MIRROR

When a logical drive is mirrored, identical data is written to a pair of physical drives, while reads are performed in parallel. The reads are performed using elevator seek and load balancing techniques where the workload is distributed in the most efficient manner. Whichever drive is not busy and is positioned closer to the data is accessed first.

With RAID 1, if one physical drive fails or has errors, the other mirrored physical drive continues to function. Moreover, if a spare physical drive is present, the spare drive is used as the replacement drive and data begins to mirrored to it from the remaining good drive

The logical drive’s data capacity equals the smaller physical drive. For example, a 100 GB physical drive and a 120 GB physical drive have a combined capacity of 100 GB in a mirrored logical drive.

If physical drives of different capacities are used, there is unused capacity on the larger drive.

RAID 1 logical drives on Vess R2600 consist of two physical drives.

If you want a mirrored logical drive with more than two physical drives, see “RAID 1E – Enhanced Mirror”.

Advantages	Disadvantages
<ul style="list-style-type: none">• Simplest RAID storage subsystem design• Can increase read performance by processing data requests in parallel since the same data resides on two different drives	<ul style="list-style-type: none">• Very high disk overhead – uses only 50% of total capacity

Recommended Applications for RAID 1:

- Accounting
- Payroll
- Financial
- Any application requiring very high availability

RAID 1E – ENHANCED MIRROR

RAID 1E offers the security of mirrored data provided by RAID 1 plus the added capacity of more than two physical drives. It also offers overall increased read/write performance plus the flexibility of using an odd number of physical drives. With RAID 1E, each data stripe is mirrored onto two physical drives. If one drive fails or has errors, the other drives continue to function, providing fault tolerance.The advantage of RAID 1E is the ability to use an odd number of physical drives, unlike RAID 1 and RAID 10. You can also create a RAID 1E Logical Drive with an even number of physical drives. However, with an even number of drives, you obtain somewhat greater security with comparable performance using RAID 10.

RAID 1E logical drives consist of three or more physical drives. You can create an array with just two physical drives and specify RAID 1E. But the resulting array is actually a RAID 1.

Advantages	Disadvantages
<ul style="list-style-type: none">• Implemented as a mirrored disk array whose segments are RAID 0 disk arrays• High I/O rates are achieved thanks to multiple stripe segments• Can use an odd number of disks	<ul style="list-style-type: none">• Very high disk overhead – uses only 50% of total capacity

Recommended Applications for RAID 1E:

- Imaging applications
- Database servers
- General fileserver

RAID 3 - BLOCK STRIPE AND DEDICATED PARITY

RAID 3 organizes block data across multiple physical drives and parity data on a dedicated drive. Generally, RAID Level 3 tends to exhibit lower random write performance due to the heavy workload of parity recalculation for each I/O. Heavy I/O loads with a large number of writes tends to work the parity drive harder in proportion to the other drives in the logical drive. Promise implements RAID 3 with block-level striping, rather than byte-level striping.

Advantages	Disadvantages
<ul style="list-style-type: none">• Good for large files with high transfer rates where fault tolerance is also required	<ul style="list-style-type: none">• Random write performance only fair. Disk failure has a medium impact on throughput

Recommended Applications for RAID 3:

- Video editing
- Multimedia production
- File and application server

RAID 5 – BLOCK AND PARITY STRIPE

RAID 5 organizes block data and parity data across the physical drives. Generally, RAID Level 5 tends to exhibit lower random write performance due to the heavy workload of parity recalculation for each I/O. RAID 5 is generally considered to be the most versatile RAID level. It works well for file, database, application and web servers.

The capacity of a RAID 5 logical drive equals the smallest physical drive times the number of physical drives, minus one. Hence, a RAID 5 logical drive with four 100 GB physical drives has a capacity of 300 GB. A RAID 5 logical drive with two 120 GB physical drives and one 100 GB physical drive has a capacity of 200 GB.

RAID 5 is generally considered to be the most versatile RAID level.

A RAID 5 on Vess R2600 consists of 3 to 32 physical drives.

Advantages	Disadvantages
<ul style="list-style-type: none">• High Read data transaction rate• Medium Write data transaction rate• Good aggregate transfer rate• Most versatile RAID level	<ul style="list-style-type: none">• Disk failure has a medium impact on throughput

Recommended Applications for RAID 5:

- File and Application servers
- WWW, E-mail, and News servers
- Intranet servers

RAID 6 – BLOCK AND DOUBLE PARITY STRIPE

RAID level 6 stores dual parity data is rotated across the physical drives along with the block data. A RAID 6 logical drive can continue to accept I/O requests when any two physical drives fail.

Hence, a RAID 6 logical drive with (7) 100 GB physical drives has a capacity of 500 GB. A RAID 6 logical drive with (4) 100 GB physical drives has a capacity of 200 GB.

RAID 6 becomes more capacity efficient in terms of physical drives as the number of physical drives increases.

RAID 6 provides double fault tolerance. Your logical drive remains available when up to two physical drives fail.

RAID 6 is generally considered to be the safest RAID level.

A RAID 6 on Vess R2600 consists of 4 to 32 physical drives.

Advantages	Disadvantages
<ul style="list-style-type: none">• High Read data transaction rate• Medium Write data transaction rate• Good aggregate transfer rate• Safest RAID level, except for RAID 60	<ul style="list-style-type: none">• High disk overhead – equivalent of two drives used for parity• Slightly lower performance than RAID 5

Recommended Applications for RAID 6:

- Accounting and Financial
- Database servers
- Any application requiring very high availability

RAID 10 – MIRROR + STRIPE

Mirror + Stripe combines both of the RAID 1 and RAID 0 logical drive types. RAID 10 can increase performance by reading and writing data in parallel or striping, and duplicating the data, or mirroring.

PROMISE implements RAID 10 by creating a data stripe over one pair of disk drives, then mirroring the stripe over a second pair of disk drives. Some applications refer to this method as RAID 0+1

The data capacity RAID 10 logical drive equals the capacity of the smallest physical drive times the number of physical drives, divided by two.

In some cases, RAID 10 offers double fault tolerance, depending on which physical drives fail.

RAID 10 arrays require an even number of physical drives and a minimum of four.

For RAID 10 characteristics using an odd number of physical drives, choose RAID 1E.

Advantages	Disadvantages
<ul style="list-style-type: none">Implemented as a mirrored disk array whose segments are RAID 0 disk arraysHigh I/O rates due to multiple stripe segments	<ul style="list-style-type: none">Very high disk overhead – uses only 50% of total capacity

Recommended Applications for RAID 10:

- Imaging applications
- Database servers
- General fileserver

RAID 30 – STRIPING OF DEDICATED PARITY

RAID 30 combines both RAID 3 and RAID 0 features. Data is striped across physical drives as in RAID 0, and it uses dedicated parity as in RAID 3. RAID 30 provides data reliability, good large file read, and high transfer rate performance.

Component	Minimum	Maximum
Number of Axles	2	16
Physical Drives per Axle	3	32
Physical Drives per Logical Drive	6	256

RAID 30 AXLES

When you create a RAID 30 you must specify the number of axles. An axle refers to two or more RAID 3 logical drives striped together to make a RAID 30. An axle can have from 3 to 32 physical drives, depending on the number of physical drives in the logical drive. The chart below shows RAID 30 logical drives with 6 to 16 physical drives, the available number of axles, and the resulting distribution of physical drives on each axle.

No. of Drives	No. of Axles	Drives per Axle		No. of Drives	No. of Axles	Drives per Axle
6	2	3,3		14	2	7,7
7	2	3,4			3	4,5,5
8	2	4,4			4	3,3,4,4
9	2	4,5		15	2	7,8
	3	3,3,3			3	5,5,5
10	2	5,5			4	3,4,4,4
	3	3,3,4			5	3,3,3,3,3
11	2	5,6		16	2	8,8
	3	3,4,4			3	5,5,6
12	2	6,6			4	4,4,4,4
	3	4,4,4			5	3,3,3,3,4
	4	3,3,3,3				
13	2	6,7				
	3	4,4,5				
	4	3,3,3,4				

Advantages	Disadvantages
<ul style="list-style-type: none">Good for large files with high transfer rates where fault tolerance is also required	<ul style="list-style-type: none">Random write performance only fair. Disk failure has a medium impact on throughput

Recommended Applications for RAID 30:

- Video editing
- Multimedia production
- File and application server

RAID 50 – STRIPING OF DISTRIBUTED PARITY

RAID 50 combines both RAID 5 and RAID 0 features. Data is striped across physical drives as in RAID 0, and it uses distributed parity as in RAID 5. RAID 50 provides data reliability, good overall performance, and supports larger volume sizes.

The data capacity RAID 50 logical drive equals the capacity of the smallest physical drive times the number of physical drives, minus two.

RAID 50 also provides very high reliability because data is still available even if multiple physical drives fail (one in each axle). The greater the number of axles, the greater the number of physical drives that can fail without

the RAID 50 logical drive going offline.

Component	Minimum	Maximum
Number of Axles	2	16
Physical Drives per Axle	3	32
Physical Drives per Logical Drive	6	256

RAID 50 AXLES

When you create a RAID 50, you must specify the number of axles. An axle refers to a single RAID 5 logical drive that is striped with other RAID 5 logical drives to make RAID 50. An axle can have from 3 to 32 physical drives, depending on the number of physical drives in the logical drive.

The chart below shows RAID 50 logical drives with 6 to 32 physical drives, the available number of axles, and the resulting distribution of physical drives on each axle.

RAID 50 Logical Drive						
No. of Drives	No. of Axles	Drives per Axle		No. of Drives	No. of Axles	Drives per Axle
6	2	3,3		14	2	7,7
7	2	3,4			3	4,5,5
8	2	4,4			4	3,3,4,4
9	2	4,5		15	2	7,8
	3	3,3,3			3	5,5,5
10	2	5,5			4	3,4,4,4
	3	3,3,4			5	3,3,3,3,3
11	2	5,6		16	2	8,8
	3	3,4,4			3	5,5,6
12	2	6,6			4	4,4,4,4
	3	4,4,4			5	3,3,3,3,4
	4	3,3,3,3				
13	2	6,7				
	3	4,4,5				
	4	3,3,3,4				

Advantages	Disadvantages
<ul style="list-style-type: none">High Read data transaction rateMedium Write data transaction rateGood aggregate transfer rateHigh reliabilitySupports large volume sizes	<ul style="list-style-type: none">Higher disk overhead than RAID 5

Recommended Applications for RAID 50:

- File and Application servers
- Transaction processing
- Office application with many users accessing small files

RAID 60 – STRIPING OF DOUBLE PARITY

RAID 60 combines both RAID 6 and RAID 0 features. Data is striped across disks as in RAID 0, and it uses double distributed parity as in RAID 6. RAID 60 provides data reliability, good overall performance and supports larger volume sizes.

The total capacity of a RAID 60 logical drive is the smallest physical drive times the number of physical drives, minus four.

RAID 60 also provides very high reliability because data is still available even if multiple physical drives fail (two in each axle). The greater the number of axles, the greater the number of physical drives that can fail without the RAID 60 logical drive going offline.

Component	Minimum	Maximum
Number of Axles	2	16
Physical Drives per Axle	4	32
Physical Drives per Logical Drive	8	256

RAID 60 AXLES

When you create a RAID 60, you must specify the number of axles. An axle refers to a single RAID 6 logical drive that is striped with other RAID 6 logical drives to make RAID 60. An axle can have from 4 to 32 physical drives, depending on the number of physical drives in the logical drive.

RAID 60 Logical Drive							
No. of Drives		No. of Axles	Drives per Axle		No. of Drives	No. of Axles	Drives per Axle
8	2	4,4		17	2	8,9	
9	2	4,5			3	5,6,6	
10	2	5,5			4	4,4,4,5	
11	2	5,6		18	2	9,9	
12	2	6,6			3	6,6,6	
	3	4,4,4			4	4,4,5,5	
13	2	6,7		19	2	9,10	
	3	4,4,5			3	6,6,7	
14	2	7,7				4	4,5,5,5
	3	4,5,5		20	2	10,10	
15	2	7,8			3	6,7,7	
	3	5,5,5			4	5,5,5,5	
16	2	8,8			5	4,4,4,4,4	
	3	5,5,6					
	4	4,4,4,4					

Advantages	Disadvantages
<ul style="list-style-type: none">• High Read data transaction rate• Medium Write data transaction rate• Good aggregate transfer rate• Safest RAID level	<ul style="list-style-type: none">• High disk overhead – equivalent of two drives used for parity• Slightly lower performance than RAID 50


- Recommended Applications for RAID 60:
- Accounting and Financial
 - Database servers
 - Any application requiring very high availability

RAID LEVEL MIGRATION

- The term “Migration” means either or both of the following:
- Change the RAID level of a logical drive.
 - Expand the storage capacity of a logical drive.

On Vess R2600, RAID level migration is performed on the disk array but it applies to the logical drives. Migration does not disturb your data. You can access the data while the migration is in progress. When migration is done, your disk array has a different RAID level and/or a larger capacity.

MIGRATION REQUIREMENTS

- The following conditions affect RAID level migration:
- The disk array and logical drive must show a green check  icon.
 - The Target disk array may require more physical drives than the Source disk array.
 - If the Target disk array requires an EVEN number of physical drives but the Source disk array has an ODD number, ADD a physical drive as part of the migration process.
 - You cannot reduce the number of physical drives in your disk array, even if the Target disk array requires fewer physical drives than the Source disk array.
 - RAID 1 (mirroring) works with two drives only. Only a single-drive RAID 0 disk array can migrate to RAID 1. Other RAID Levels use too many drives to migrate.
 - You cannot migrate a disk array when it is Critical or performing activities such as Synchronizing, Rebuilding, and PDM.
 - For RAID 6 or RAID 60, you can only migrate between these two RAID levels. Destination RAID 60 axles can have up to 16 physical drives. Other limitations might apply.

SOURCE AND TARGET RAID LEVELS

The tables on the following pages show the migration options for each source logical drive by its RAID level. The available target RAID levels are shown with their requirements.

RAID 0

A RAID 0 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 0	Add physical drives.
RAID 1	2 physical drives only. Only a single-drive RAID 0 can migrate to RAID 1 by adding 1 physical drive.
RAID 1E	3 or more physical drives. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 5	3 physical drives minimum, 32 maximum. RAID 0 must have less than 16 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 6	4 physical drives minimum, 32 maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	4 physical drives minimum. Even number of physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 1

A RAID 1 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 0	None.
RAID 1E	3 or more physical drives. Add 1 or more physical drives.
RAID 5	3 physical drives minimum, 32 maximum. RAID 1 must have less than 32 physical drives. Add 1 or more physical drives.
RAID 10	4 physical drives minimum. Even number of physical drives. Add 2 or more physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum. Add 4 or more physical drives.

RAID 1E

A RAID 1E Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 0	None.
RAID 1E	Add physical drives.
RAID 5	3 physical drives minimum, 32 maximum. RAID 1E must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 10	4 physical drives minimum. Even number of physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum.

RAID 3

A RAID 3 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 0	None.
RAID 1E	None
RAID 3	Add physical drives. 32 maximum.
RAID 5	3 physical drives minimum. 32 maximum. If existing physical drives have no unused space add 1 or more physical drives..
RAID 6	4 physical drives minimum, 32 maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	4 physical drives minimum. Even number of physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 30, 50	6 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 5

A RAID 5 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 0	None.
RAID 1E	None.
RAID 5	Add physical drives. 32 maximum.
RAID 6	4 physical drives minimum, 32 maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	4 physical drives minimum. Even number of physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.
----------------	---

RAID 6

A RAID 6 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 6	Add physical drives. 32 maximum.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 10

A RAID 10 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 0	None.
RAID 1E	None.
RAID 5	3 physical drives minimum, 32 maximum. RAID 10 must have less than 16 physical drives.
RAID 6	4 physical drives minimum, 32 maximum. RAID 10 must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	Add physical drives. Even number of physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

When you migrate RAID 10 logical drive, it becomes RAID 1E by default.

If you want a RAID 10 logical drive, there must be an even number of physical drives and you must specify RAID 10 for the target logical drive.

RAID 30

A RAID 30 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 0	None.
RAID 1E	None.

RAID 3, 5	32 physical drives maximum. RAID 30 must have less than 32 physical drives.
RAID 6	32 physical drives maximum. RAID 30 must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	Even number of physical drives.
RAID 30	Add physical drives. 32 per axle maximum.
RAID 50	6 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 50

A RAID 50 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 0	None.
RAID 1E	None.
RAID 5	32 physical drives maximum. RAID 50 must have less than 32 physical drives.
RAID 6	32 physical drives maximum. RAID 50 must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	Even number of physical drives.
RAID 50	Add physical drives. 32 per axle maximum.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

You can add physical drives to a RAID 50 array but you cannot change the number of axles.

RAID 60

A RAID 60 Source logical drive can migrate to the following Target logical drives:

Target	Requirements
RAID 6	32 physical drives maximum. RAID 60 must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 60	Add physical drives. 32 per axle maximum.

You can add physical drives to a RAID 60 array but you cannot change the number of axles.

Stripe Size

Stripe Size, also called “Stripe Block Size,” refers to the size of the data blocks written to, and read from, the physical drives. Stripe Size is specified when you create a logical drive. You can choose Stripe Size directly when you use the Wizard Advanced Configuration function to create a logical drive.

You cannot change the Stripe Size of an existing logical drive. You must delete the logical drive and create a new one.

The available Stripe Sizes are 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB. 64 KB is the default. There are two issues to consider when choosing the Stripe Size:

- You should choose a Stripe Size equal to, or smaller than, the smallest cache buffer found on any physical drive in the disk array. Selecting a larger value slows read/write performance because physical drives with smaller cache buffers need more time for multiple accesses to fill their buffers.
- If your data retrieval consists of fixed data blocks, such as with some database or video applications, then you should choose that size as your Stripe Size.

If you do not know the cache buffer or fixed data block sizes, choose 64 KB as your Stripe Size. Generally speaking,

- Email, POS, and web servers prefer smaller stripe sizes.
- Video and database applications prefer larger stripe sizes.

Sector Size

A sector is the smallest addressable area on a physical drive. Sector size refers to the number of data bytes a sector can hold. A smaller sector size is a more efficient use of a physical drive’s capacity. 512 bytes (512 B) is the most common sector size, and the default in WebPAM PROe.

Preferred Controller ID

When you create a logical drive using the Advanced method of disk array creation, you can specify the Preferred Controller ID:

- **Controller 1** – Assign all logical drives to Controller 1
- **Controller 2** – Assign all logical drives to Controller 2.
- **Automatic** – Alternate logical drive assignments between Controllers 1 and 2.

Automatic is the default and preferred setting because it balances the logical drive assignments for you.

Initialization

Initialization is done to logical drives after they are created from a disk array. Full initialization sets all data bits in the logical drive to a specified pattern, such as all zeros. The action is useful because there may be residual data on the logical drives left behind from earlier configurations. For this reason, Initialization is recommended for all new logical drives. "Initializing a Logical Drive (CLU)" on page 140.



Caution

When you initialize a logical drive, all the data on the logical drive is lost. Backup any important data before you initialize a logical drive.

Partition and Format

Like any other type of fixed disk media in your system, a RAID logical drive must also be partitioned and formatted before use. Use the same method of partitioning and formatting on an logical drive as you would any other fixed disk.

Depending on the operating system you use, there may or may not be various capacity limitations applicable for the different types of partitions.

SPARE DRIVES

Spare drive technology includes:

- "Definition" on page 193
- "Options" on page 193
- "Requirements" on page 193
- "Transition" on page 193

DEFINITION

A spare drive is a physical drive that you designate to automatically replace the failed physical drive in a disk array.

The general recommendation is to:

- Provide at least one spare drive for every 16 physical drives in the RAID system
- Configure the spares as global revertible spare drives

OPTIONS

There are several options you can specify for a spare drive:

- System Options
 - **Revertible** – Returns to its spare drive assignment after you replace the failed physical drive in the disk array and run the Transition function.
 - **Media Patrol** – By default, Media Patrol runs on spare drives unless you disable it.
- Spare Type
 - **Global** – Can be used by any disk array
 - **Dedicated** – Can be used only by the assigned disk array
- Media Type (type of physical drive)
 - Hard Disk Drive (HDD)
 - Solid State Drive (SSD)

REQUIREMENTS

The spare drive must:

- Have adequate capacity to replace the largest physical drive in your disk arrays.
- Be the same media type as the physical drives in your disk arrays.

A revertible spare drive requires:

- You to replace the failed physical drive in the disk array
- You to run the Transition function

TRANSITION

Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare. The revertible spare drive returns to its original status. In order to run the Transition function, the spare drive must be revertible.

In addition, you must specify an unconfigured physical drive of the same or larger capacity and same media type as the revertible spare drive.

RUNNING A TRANSITION

The Transition feature enables you to specify “permanent” spare drives for your Vess R2600 subsystem. Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare. The revertible spare drive returns to its original status.

Transition happens automatically when the following sequence of events takes place:

- You create a revertible spare drive. See “Creating a Spare Drive Manually” on page 102.
- A physical drive assigned to your disk array fails and the array goes critical or degraded.
- Vess R2600 automatically rebuilds your array to the revertible spare drive and the array becomes functional again.
- You replace the failed physical drive with a new physical drive of equal or greater capacity.
- Vess R2600 automatically transitions (moves) the data from the revertible spare to the new physical drive.
- The new physical drive becomes part of the array and the revertible spare drive returns to its original spare status.

Transition happens manually when you specify a different unconfigured physical drive to transition (move) the data from the revertible spare drive.

See the example on the following pages.

EXAMPLE

Following is an example to explain the Transition function.

In the example above, there is a four-drive RAID 5 disk array and a global spare drive. Physical drives 1, 2, 3, and 4 belong to the disk array. Physical drive 5 remains unconfigured. Physical drive 6 is a revertible spare drive.

If a physical drive fails in a disk array and there is a spare drive of adequate capacity available, the controller automatically rebuilds the array using the spare drive. In this example, physical drive 3 failed and the array is rebuilt using physical drive 6, the revertible spare drive.

When the rebuild is complete, the spare drive has replaced the failed drive. In this example, failed drive 3 was replaced by spare drive 6. The disk array now consists of physical drives 1, 2, 4, and 6.

There is no spare drive at this moment. Even if physical drive 5 is of adequate capacity, it has not been designated as a spare, therefore the controller cannot use it as a spare.

AUTOMATIC TRANSITION

At this juncture, you would replace the failed drive in slot 3 with a new one of the same or greater capacity.

When the Vess R2600 controller detects the new drive in slot 3, the controller:

- Automatically transitions the data on drive 6 to drive 3
- Returns drive 6 to spare status

When the Automatic Transition is finished, physical drives 1, 2, 3, and 4 belong to the disk array and physical drive 6 is a revertible spare drive. The original configuration is restored.

MANUAL TRANSITION

If you wanted to use the drive in slot 5 as a member of the disk array, rather than the drive in slot 3, you would run the Transition function manually.

When the Manual Transition is finished, physical drives 1, 2, 4, and 5 belong to the disk array and physical drive 6 is a revertible spare drive.

At this point, you would replace the drive in slot 3. The new drive in slot 3 remains unconfigured until you assign it to a disk array or as a spare.

RAID CONTROLLERS

RAID controller technology includes;

- "LUN Affinity" on page 194
- "ALUA" on page 194
- "Cache Policy" on page 195
- "Preferred Controller ID" on page 196
- "Power Saving" on page 196
- "Capacity Coercion" on page 196

LUN AFFINITY

Vess R2600 subsystems with dual RAID controllers include a LUN Affinity feature. Normally, either controller can access all logical drives. LUN Affinity enables you to specify which controller can access each logical drive. Use this feature to balance the load of your logical drives between the two controllers.

To use LUN Affinity you must:

- Have two RAID controllers in the subsystem.
- Set the redundancy type to Active-Active.
 - See "Making Subsystem Settings (CLU)" on page 122.
- Enable LUN Affinity.
 - See "Making Subsystem Settings (CLU)" on page 122..

On subsystems with two RAID controllers, when Cache Mirroring is disabled, LUN Affinity is enabled automatically.

ALUA

Vess R2600 supports Asymmetric Logical Unit Access (ALUA) on Linux OSes. ALUA is a multithipathing tool. It enables an initiator (your host PC or server) to discover target port groups that provide a common failover / failback behavior for your LUNs. ALUA enables the host to see which paths are in an optimal state and which are not.

To use ALUA you must:

- Have two RAID controllers in the subsystem.

- Set the redundancy type to Active-Active.
See "Making Subsystem Settings (CLU)" on page 122.
- Enable LUN Affinity and ALUA.
See "Making Controller Settings (CLU)" on page 124.

For more information, see "Appendix C: Multipathing on Linux" on page 246.

CACHE POLICY

As it is used with Vess R2600, the term cache refers to any of several kinds of high-speed, volatile memory that hold data moving from your computer to the physical drives or vice-versa. Cache is important because it can read and write data much faster than a physical drive. There are read caches, which hold data as it is read from a physical drive; and write caches, which hold data as it is written to a physical drive.

In order to tune the cache for best performance in different applications, user-adjustable settings are provided. Cache settings are made on the RAID controller. See "Making Controller Settings (CLU)" on page 124.

READ CACHE POLICY

- **Read Cache** – The read cache is enabled but no pre-fetch action.
- **Read Ahead** – The read cache and predictive pre-fetch feature are enabled. Read-ahead anticipates the next read and performs it before the request is made. Can increase read performance.
- **Forced Read Ahead** – The read cache and aggressive pre-fetch feature are enabled. See “Forced Read-Ahead Cache” below.
- No Cache – The read cache is disabled.

WRITE CACHE POLICY

- **Write Back** – Data is written first to the cache, then to the logical drive. Better performance. Vess R2600 has a cache backup battery to protect data in the cache from a sudden power failure.
- **Adaptive Writeback** – See “Adaptive Writeback Cache” below.
- **Write Thru** – Also “Write Through.” Data is written to the cache and the logical drive at the same time. Safer. If your write cache policy is set to Write Back, the write policy automatically changes to Write Thru when all of the following conditions occur:
 - The logical drive write policy is set to Write Back
 - The Adaptive Writeback Cache feature is enabled

- The cache backup battery goes offline

When the battery comes back online, the write policy automatically changes back to Write Back.

FORCED READ-AHEAD CACHE

On the Vess R2600 subsystem, you can set the logical drive read cache policy to Forced Read Ahead an enable the aggressive pre-fetch feature.

The Forced Read-Ahead cache policy setting provides predictive pre-fetching of data requests, allowing the controller to aggressively buffer large chunks of data in cache memory to prevent frame drops on high-bandwidth video playback. Not normally enabled for non-video applications.

ADAPTIVE WRITEBACK CACHE

On the Vess R2600 subsystem, you can set the logical drive write cache policy to Write Thru or Write Back. If you set the write cache policy to Write Back, your data is first written to the controller cache, and later to the logical drive. This action improves performance. To preserve the data in the cache in the event of a power failure, the subsystem has a backup battery that powers the cache. To see an estimate of how long the battery can power the cache, see “Viewing Battery Information” on page 89 or “Viewing Battery Information” on page 222.

The Adaptive Writeback Cache feature protects your data by changing the write cache settings while the cache backup battery is offline. When all of the following conditions occur:

- The logical drive write policy is set to Write Back.
- The Adaptive Writeback Cache feature is enabled.
- The cache backup battery goes offline.

The write policy automatically changes to Write Thru. When the battery comes back online, the write policy automatically changes back to Write Back.

To enable the Adaptive Writeback Cache option, see "Making Controller Settings (CLU)" on page 124.

HOST CACHE FLUSHING

On the Vess R2600 subsystem, you can enable or disable host cache flushing.

When enabled, host cache flushing guards against data loss in the event of a power failure. However RAID performance is slightly reduced.

When disabled, the Vess R2600 subsystem has greater sustained bandwidth and lower latency, which are helpful for real-time video capture.

When you operate the Vess R2600 with host cache flushing disabled, use a UPS to protect against data loss.

ADVANCED BATTERY FLASH BACKUP

This feature uses power from the BBU to save the write cache used for the RAID controller to NAND flash memory. When the write chache data backup is complete, the remaining BBU ceases to supply power for the write cache and the remaining battery capacity is conserved. This in effect, indefinitely extends the period of write loss protection beyond the required 72 hours typical for most BBU setups.

PREFERRED CONTROLLER ID

SeePreferred Controller ID3.

POWER SAVING

Power saving is a method of conserving energy by applying specific actions to hard disk drives (HDD). After an HDD has been idle for the set period of time, you can elect to:

- Park the read/write heads – Referred to as **Power Saving Idle Time** on Vess R2600.
- Reduce disk rotation speed – Referred to as **Power Saving Standby Time** on Vess R2600.
- Spin down the disk (stop rotation) – Referred to as **Power Saving Stopped Time** on Vess R2600.

Power management must be:

- Set on the RAID controller. see "Making Controller Settings (CLU)" on page 124.
- Enabled on each HDD.

CAPACITY COERCION

This feature is designed for fault-tolerant logical drives (RAID 1, 1E, 5, 10, 50, and 60). It is generally recommended to use physical drives of the same size in your disk arrays. When this is not possible, the system adjusts for the size differences by reducing or coercing the capacity of the larger drives to match the smaller ones. With Vess R2600, you can choose to enable capacity coercion and any one of four methods.

Enable capacity coercion and choose a method, see "Making Controller Settings (CLU)" on page 124. The choices are:

- GB Truncate – (Default) Reduces the useful capacity to the nearest 1,000,000,000 byte boundary.
- 10GB Truncate – Reduces the useful capacity to the nearest 10,000,000,000 byte boundary.
- Group Rounding – Uses an algorithm to determine how much to truncate. Results in the maximum amount

of usable drive capacity.

- **Table Rounding** – Applies a predefined table to determine how much to truncate.

Capacity coercion also affects a replacement drive used in a disk array. Normally, when an physical drive fails, the replacement drive must be the same capacity or larger. However, the capacity coercion feature permits the installation of a replacement drive that is slightly smaller (within 1 gigabyte) than the remaining working drive. For example, the remaining working drives can be 80.5 GB and the replacement drive can be 80.3, since all are rounded down to 80 GB. This permits the smaller drive to be used.

Without capacity coercion, the controller does not permit the use of a replacement physical drive that is slightly smaller than the remaining working drives.

iSCSI MANAGEMENT

iSCSI management uses the following terms:

- "iSCSI on a VLAN" on page 197
- "Initiator" on page 198
- "Target" on page 198
- "Portal" on page 198
- "Port" on page 199
- "Trunk" on page 199
- "Session" on page 199
- "iSNS" on page 199
- "CHAP" on page 199
- "Ping" on page 199

Also see "Managing iSCSI Connections" on page 109.

A detailed explanation of iSCSI functions and how to best use them is beyond the scope of this document. For more information, contact the Internet Engineering Task Force at <http://www.ietf.org/>.

iSCSI ON A VLAN

Vess R2600 supports up to 32 iSCSI portals. Each iSCSI portal can belong to a different VLAN for a maximum of 32 VLANs.

To set up the Vess R2600 subsystem for a VLAN:

1. Add a new portal with a **VLAN association**.

See "Adding iSCSI Portals" on page 111.

Note which iSCSI port you chose for the portal.

2. Connect your iSCSI data cable to the iSCSI port you chose for the new portal.

See "iSCSI Storage Area Network (SAN)" on page 24.

3. Add your iSCSI initiators to the Vess R2600's initiator list.

See "Adding an iSCSI Initiator" on page 105.

For information, see:

- "Managing iSCSI Connections" on page 109.

INITIATOR

An initiator functions as the client, in this case, your host PC or server. The initiator makes requests to and receives responses from an iSCSI target on the Vess R2600 RAID subsystem.

Each initiator has a unique iSCSI qualified name (IQN). You specify the initiator by that name when you map a LUN or logical drive to the initiator. Initiators come in two varieties, software and hardware.

SOFTWARE

A software initiator uses code to implement iSCSI. The software emulates SCSI devices for a computer by speaking the iSCSI protocol. Software initiators are available for most mainstream operating systems, and this type is the most common mode of deploying iSCSI on computers.

For more information, see your iSCSI driver user documentation.

HARDWARE

A hardware initiator uses dedicated hardware in combination with software running on it, to implement iSCSI. A common example is an iSCSI host bus adapter (HBA) card.

The iSCSI HBA is a 1-gigabit or 10 gigabit Ethernet network Interface card (NIC) that plugs into a PCI-Express slot. It looks like a SCSI device to the host PC or server’s operating system.

The iSCSI HBA uses a TCP/IP Offload Engine (TOE) to perform iSCSI and TCP processing and managing interrupts, leaving the host PC or server’s microprocessor free to run other applications.

For more information, see your iSCSI HBA user documentation.

TARGET

The target represents a storage device, in this case the Vess R2600 RAID subsystem. Each target has a unique iSCSI qualified name (IQN).

Vess R2600 supports one iSCSI target. A maximum of 256 logical drives can be mapped to a target.

Target options include Digests, CHAPs and Nop-In.

DIGESTS

A **header** digest adds a 32-bit CRC digest to detect data corruption in the header portion of each iSCSI packet.

A **data** digest adds a 32-bit CRC digest to detect data corruption in the data portion of each iSCSI packet.

If a data packet arrives with an invalid CRC digest, the data packet is rejected.

Header and data digests work best with initiators equipped with a TOE. Refer to your iSCSI HBA. For more information, see your iSCSI HBA user documentation.

CHAPs

Challenge Handshake Authentication Protocol (CHAP) is an authentication mechanism used to authenticate iSCSI sessions between initiators and targets.

A uni-directional or peer CHAP authenticates from the target (Vess R2600) to the initiator (host PC or server).

A bi-directional or local CHAP authenticates target to initiator and initiator to target.

NOP-IN

Nop-In is used to check iSCSI connection status on the specified target. Nop-In periodically sends a PDU to the initiator that requests a response. If there is no reply, the Vess R2600 will disconnect the iSCSI connection in order to preserve system resources for other operations. The default setting is enabled.

PORTAL

A portal is the logical point of connection between the Vess R2600 and the iSCSI network. Portals use an IP address and a TCP port number to identify an IP storage resource. Vess R2600 supports up to 32 iSCSI portals. Vess R2600 uses TCP port 3260.

Vess R2600 supports both IPv4 and IPv6 addresses. See “Internet Protocols” on page 373.

Portals on Vess R2600 support three types of port associations:

- PHY – A simple connection through one port.
- VLAN – Virtual Local Area Network. The portal is part of a virtual network. Used when a dedicated network is not available for iSCSI.
- Trunk – An aggregation of two or more iSCSI ports on the same RAID controller. Also known as a link aggregation. This feature combines ports to increase bandwidth.

Once you have made a port association, you cannot change it. If you have no portals with the port association you want, create a new portal.

Each iSCSI portal can belong to a different VLAN. Vess R2600 supports 32 VLANs.

PORT

A port is the physical point of connection between the Vess R2600 and the iSCSI network. There are four ports on each RAID controller for a total of eight on a system with dual controllers. When you create a portal, you specify one port.

There are two options for each iSCSI port:

- Jumbo Frame – Enables jumbo frame support on the port.

The standard Ethernet frame is 1518 bytes, with 1500 bytes for payload. A jumbo frame ranges from 1500 bytes to 9000 bytes of payload. Because jumbo frames carry more data, they are used to reduce network management overhead, thereby increasing network throughput.

TRUNK

A trunk is an aggregation of two or more iSCSI ports on the same RAID controller. Also known as a **link aggregation**. This feature combines ports to increase bandwidth. Ports must be **enabled** to add them to a trunk. Trunks are identified by their Trunk IDs.

When you create a trunk, you specify:

- Controller ID – RAID controller whose iSCSI ports you are using.
- Master port – Any available iSCSI port.
- Slave ports – The remaining available iSCSI ports.

SESSION

A session is a group of TCP connections that link an iSCSI initiator with a target. Each RAID controller supports a maximum of 128 sessions, or 256 per subsystem.

iSNS

Internet Storage Name Service (iSNS) is a protocol that facilitates automated discovery, management, and configuration of iSCSI devices on a TCP/IP network. iSNS service runs on an iSNS server on your network. You can enable iSNS on the Vess R2600 and specify the IP address and port number of the iSNS server.

CHAP

Challenge Handshake Authentication Protocol (CHAP) is an authentication mechanism used to authenticate iSCSI sessions between initiators and targets. A uni-directional or peer CHAP authenticates from the target (Vess R2600) to the initiator (host PC or server). A bi-directional or local CHAP authenticates target to initiator and initiator to target .

PING

Ping is a computer network administration utility that tests whether a device is accessible over the IP network. Ping sends echo request packets to the target node, such as your host PC or server, and waits for a response. It measures the time from transmission to reception and records any packet loss. Vess R2600 can ping through its virtual management port and each of its iSCSI data ports. You must input the IP address of the target client.

INTERNET PROTOCOLS

Vess R2600 supports the IPv4 and IPv6 protocols.

Protocol	Addresses		Example
IPv4	32-bits	4.3 x 10 ⁹	10.0.0.1
IPv6	128-bits	3.4 x 10 ³⁸	2001:0000:0000:0000:0000:0000:e2a8:4337 Abbreviated 2001:0:0:0:0:0:e2a8:4337

CHAPTER 8: TROUBLESHOOTING

This chapter contains the following topics:

- Vess R2600 is Beeping (below)
- LEDs Display Amber or Red (page 202)
- CLU Reports a Problem (page 206)
- WebPAM PROe Reports a Problem (page 208)
- USB Support Reports a Problem (page 209)
- "Enclosure Problems" on page 209
- "RAID Controller Problems" on page 211
- "Disk Array and Logical Drive Problems" on page 213
- "Physical Drive Problems" on page 213
- "Connection Problems" on page 216
- "Power Cycling the Subsystem" on page 218
- "Event Notification Response" on page 218

Vess R2600 is BEEPING

Vess R2600’s alarm has five different patterns.

When you first power-up the Vess R2600, it beeps twice to show normal operation.

The audible alarm sounds at other times to inform you that the Vess R2600 needs attention. But the alarm does not specify the condition.

When the alarm sounds:

- Check the front and back of Vess R2600 enclosure for red or amber LEDs.
- If email notification is enabled, check for new messages.
- Check for yellow ! 🟡 red X 🔴 icons.
- Check the event log.

See "Viewing Runtime Events" on page 206 and "Viewing NVRAM Events" on page 207.

When a continuous tone sounds, there are multiple alarm patterns sounding at the same time.

SILENCING THE BUZZER



Caution

This action disables the buzzer for all events.

To silence the buzzer:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the **Buzzer** and click the **Settings** button.
4. Uncheck the **Enable Buzzer** box.
5. Click the **Save** button.

LEDs DISPLAY AMBER OR RED

LEDs are used on Vess R2600’s:

- Front Panel LEDs (page 202)
- Drive Carrier LEDs (page 203)
- Back Panel LEDs (page 203)

FRONT PANEL LEDs

When the power is switched on, the LEDs on the front of the Vess R2600 light up.

When boot-up is finished and the Vess R2600 is functioning normally:

- Power, Global Enclosure Status, and Global RAID Status LEDs display green continuously.
- Controller Activity LED flashes green when there is controller activity.
- System Heartbeat LED blinks blue (once a second), and repeats the pattern.


Steady means the LED is on.

Blinking means a regular on/off pattern.

Flashing means an intermittent and irregular on/off pattern.

Dark means the LED is off.

Figure 8 - 1: Front panel LEDs right side



See "Enclosure Problems" on page 209, "RAID Controller Problems" on page 211, and "Disk Array and Logical Drive Problems" on page 213 for more information.

The Locator feature is triggered from WebPAM PROe or the CLU. It causes the LEDs to blink on and off for one minute. That action helps you find the physical component.

LED State	Power	Global Enclosure Status	Global RAID Status	Controller Activity	System Heartbeat
Dark	No power	—		No Controller in Slot	—
Steady Green		All devices normal	All LDs are on line		—
Steady Blue	Normal			No activity	
Blinking Blue	—	—	—	Activity	Normal**
Flashing Green	—	Locating device	—		—
Amber	—	One or two devices in error	One or more LD is critical; none are offline	—	—
Red	—	Three or more devices in error	One or more LD is offline	—	—

DRIVE CARRIER LEDs

The Vess R2600 spins up the disk drives sequentially to equalize power draw during start-up. After a few moments:

- The Power/Activity LED displays blue when a physical drive is present.
- The Drive Status LED displays green when the physical drive is configured as a member of a disk array or as a spare. When the physical drive is unconfigured, the LED is dark.

See the diagram and table on the next page.



Drive Carrier LEDs		
State	Power / Activity	Drive Status
Dark	No drive in carrier	Drive is unconfigured
Steady Blue	Drive is present	—
Flashing Blue	Activity on drive	—
Steady green	—	Drive is configured
Blinking	both LEDs blink	Locator feature
Amber	—	Drive is rebuilding
Red	—	Drive error or failure
* Configured means the physical drive either belongs to an array or it is assigned as a spare drive.		
* Steady means the LED is on.		
* Blinking means a regular on/off pattern.		
* Flashing means intermittent and irregular on/off pattern.		

See "Physical Drive Problems" on page 213 for a discussion of rebuilding and failed physical drives for more information.

The Locator feature is triggered from WebPAM PROe or the CLU. It causes the LEDs to blink on and off for one minute. That action helps you find the specific drive.


BACK PANEL LEDs

These LEDs give the status of the field replaceable units:

- RAID controller
- Power supply
- BBU (located on each Cooling Unit)
- Cooling Unit

These LEDs will light green to indicate normal operation. A red or amber LED indicates a problem or unit failure.

Figure 8 - 3: Power supply, BBU and Fan status LEDs



State	Battery	Power Supply	Fan
Dark	No power, Failed or not installed	No power	No power
Steady Green	Normal	Normal	Normal
Blinking Green	—	Locator feature	—
Steady Amber	—	—	Fan problem
Flashing Amber	—	—	—
Steady Red	Failed or battery removed	Failed	Failure
Flashing Red	—	—	—

CONTROLLER LED BEHAVIOR

When boot-up is finished and the Vess R2600 subsystem is functioning normally, the Controller status LED displays green continuously; the Management port LEDs display green or flash depending on your network connection; the FC, iSCSI, and SAS Expansion LEDs display green or flash during port activity.

LED	Description
SAS Expansion	Lights green when connected, flashes green when active.
Controller Status	This displays the current operational status of the controller. A steady (unblinking) green light indicates the controller is operational.
Dirty Cache	Blinks amber if cache is dirty, meaning that the controller memory cache contains data, otherwise this is dark.
USB	A steady green light indicates a valid USB connection, this is dark when not connected (no device attached).
iSCSI (2 above each port)	Left LED lights green when connected, flashes green when active, dark if not connected. Right LED indicates connection speed, green is 100 Mbps, amber is 1000 Mbps.
FC ports	See next page

Figure 8 - 4: Vess R2600i controller LEDs




Figure 8 - 5: Vess R2600fi controller LEDs




Figure 8 - 6: JBOD controller LEDs



Under normal conditions, the power supply status LEDs display green.

See "Enclosure Problems" on page 209 and "RAID Controller Problems" on page 211 for more information.

The Locator feature is triggered from WebPAM PROe or the CLU. It causes the LEDs to blink on and off for one minute. That action helps you find the physical component.

LED	Description		
FC ports (one LED above each port)	Green	Amber	Status
	dark	dark	Wake-up failure (dead board)
	off	on	POST failure (dead board)
	off	blinking slowly	Wake-up failure monitor
	off	blinking rapidly	Failure to POST
	off	flashing	POST in progress
	on	off	Failure while functioning
	on	on	Failure while functioning
	on	2 rapid blinks	Normal, link up, 2 Gb/s
	on	3 rapid blinks	Normal, link up, 4 Gb/s
	on	4 rapid blinks	Normal, link up, 8 Gb/s
	blinking slowly	off	Normal, link down
	blinking slowly	blinking slowly	Offoffline for download
	blinking slowly	blinking rapidly	Restricted offoffline mode (waiting for restart)
	blinking slowly	flashing	Restricted offoffline mode, test active

See "Enclosure Problems" on page 209 and "RAID Controller Problems" on page 211 for more information.

The Locator feature is triggered from WebPAM PROe or the CLU. It causes the LEDs to blink on and off for one minute. That action helps you find the physical component.

Checking Component Installation

To check a component’s installation, remove the component, then reinstall the component in its original location. In most cases, this action fixes a bad connection and allows Vess R2600 to detect the component. If this action does not correct the problem, replace the unit.

On Vess R2600s with dual controllers, when one controller’s Status LED is amber and the other controller’s Status LED is flashing red, it means that the controller with the flashing red LED has entered maintenance mode. See "RAID Controller Problems" on page 211.

If the Controller Status LED continues to display amber after startup, contact PROMISE Technical Support. See "Contacting Technical Support" on page 228.

The Dirty Cache LED flashes during input/output operation. If the LED shines amber and the power is off, there is unsaved data in the cache. Do NOT power down the Vess R2600 while this LED is on.

CLU Reports a Problem

The CLU reports information passively, that is you must determine which functions to check based on the sound of the Vess R2600’s audible alarm and any amber or red LEDs. SeeVess R2600 is Beeping5 andLEDs Display Amber or Red7 for more information.

Check the event logs first. Then check the reported component.

Viewing Runtime Events

To display Runtime Events:

1. From the Main Menu, highlight Event Viewer and press **Enter**.

The log of Runtime Events appears. Events are added to the top of the list. Each item includes:

- **Sequence number** – Begins with 0 at system startup.
- **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
- **Severity** – See the Table below.
- **Timestamp** – Date and time the event happened.
- **Description** – A description of the event in plain language.

2. Press the up and down arrow keys to scroll through the log.

Event Severity Levels	
Level	Description
<i>Fatal</i>	Non-recoverable error or failure has occurred.
<i>Critical</i>	Action is needed now and the implications of the condition are serious.
<i>Major</i>	Action is needed now.
<i>Minor</i>	Action is needed but the condition is not a serious at this time.
<i>Warning</i>	User can decide whether or not action is required.
<i>Information</i>	Information only, no action is required.

VIEWING NVRAM EVENTS

This screen displays a list of and information about 63 most important events over multiple subsystem startups.

To display NVRAM events:

1. From the Main Menu, highlight **Event Viewer** and press **Enter**.
2. Highlight **NVRAM Events** and press **Enter**.

The log of NVRAM Events appears. Events are added to the top of the list. Each item includes:

- **Sequence number** – Begins with 0 at system startup.
- **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
- **Severity** – See the Table on the previous page.
- **Timestamp** – Date and time the event happened.
- **Description** – A description of the event in plain language.

3. Press the up and down arrow keys to scroll through the log.

CHECKING A REPORTED COMPONENT

In this example, let us check disk array status.

1. Open the CLU.
2. Highlight **Disk Array Management** and press **Enter**.
3. Observe the status of your disk arrays.

DaId	Alias	OpStatus	CfgCapacity	FreeCapacity	MaxContiguousCap

0	DA0	OK	75.44GB	66.06GB	66.06GB
1	DA1	Degraded	189.06GB	179.68GB	179.68GB
2	DA2	OK	73.57GB	64.20GB	64.20GB

At this point, you can highlight the Degraded array and press **Enter** to see more information. See below.

Disk Array ID: 1
OperationalStatus: Degraded
FreeCapacity: 179.68 GB
SupportedRAIDLevels: 0 5 10 1E

Physical Capacity: 189.06GB
MaxContiguousCapacity: 11.18GB
ConfigurableCapacity: 179.68GB

Disk Array Alias : DA1
MediaPatrol : Enabled
PDM : Enabled

Transport
Rebuild
Predictive Data Migration
Transition
Dedicated Spare Drives in the Array
Physical Drives in the Array
Logical Drives in the Array
[Locate Disk Array]

Save Settings [CTRL-A]
Restore Settings [CTRL-R]
Return to Previous Menu

From this screen:


- Highlight **Physical Drives** in the Array and press **Enter** to identify the failed disk drive
- Highlight **Rebuild** and press **Enter** to rebuild the array after you replace the failed disk drive

WebPAM PROe Reports A Problem

WebPAM PROe reports these conditions in the header and all four tabs.

Header

The Header displays popup messages, per your configuration.



Dashboard Tab

- System Status
- Event Information

Event Severity Levels	
Level	Description
<i>Fatal</i>	Non-recoverable error or failure has occurred.
<i>Critical</i>	Action is needed now and the implications of the condition are serious.
<i>Major</i>	Action is needed now.
<i>Minor</i>	Action is needed but the condition is not a serious at this time.
<i>Warning</i>	User can decide whether or not action is required.
<i>Information</i>	Information only, no action is required.

Device Tab

- Front View, showing the drive carrier icons.
- Rear View, with Show Internal Components option.
- Physical Drive View, physical drive shown dead or offline and marked with a red X 🛑 icon.

Storage Tab

- Disk Arrays
- Logical Drives

Administration Tab

Events icon.

Event Severity Levels	
Level	Description
<i>Fatal</i>	Non-recoverable error or failure has occurred.
<i>Critical</i>	Action is needed now and the implications of the condition are serious.
<i>Major</i>	Action is needed now.
<i>Minor</i>	Action is needed but the condition is not a serious at this time.
<i>Warning</i>	User can decide whether or not action is required.
<i>Information</i>	Information only, no action is required.

Also see these troubleshooting topics:

- "Event Notification Response" on page 218
- “Enclosure Problems” on page 209.
- “Frequently Asked Questions” on page 225.

USB SUPPORT REPORTS A PROBLEM

This procedure requires a USB flash device:

- Formatted to FAT 32
- At least 50 MB of free space



Caution

Verify that there is no firmware image file on the USB flash device. If a firmware image file is present, the RAID controller might attempt a firmware update.

To collect a service report using the USB Support feature:

1. Insert the USB flash device into one of the USB ports on the front panel, left side of the enclosure. The OPAS LED blinks green in one-second intervals.
2. Wait until the OPAS LED stops blinking green and displays steady green.
3. Remove the USB flash device.
4. Insert the USB flash device into a USB port on your PC.
5. On the USB flash device, open the OPAX_XXXXXX folder to obtain the report and log.

ENCLOSURE PROBLEMS

Enclosure Problems include:

- "Diagnosing an Enclosure Problem" on page 209
- "Power Supplies" on page 210
- "Batteries" on page 210

DIAGNOSING AN ENCLOSURE PROBLEM

Check System Status on the Dashboard tab. If a yellow ! or red X appears in the System Status box:

1. Click the name link of the component with the red X icon.

The Components List of the Device tab displays.

2. Mouse-over Enclosure with the red X icon and click the View button.

The components list expands and shows the power supply fans, which server as the Cooling Unit of the Vess R2600 enclosure.

Note that the fans for power supply 2 (PSU 2) have failed.

3. Click the Back View icon on the Device tab.
4. Click the picture of the enclosure.

A popup messages displays the status of each component.

When a power supply fan fails, you must replace the power supply. See "Replacing a Power Supply" on page 177 for more information.

If the system reports a fan malfunction, contact Technical Support (see page 228) immediately to schedule replacement of the suspect power supply as soon as possible. Running the unit in this condition for more than three weeks may shorten subsystem life and void your warranty.

OVERHEATING

Overheating is a potentially serious condition because the excessively high temperatures can lead to physical drive failure and controller malfunction.

Overheating usually results from inadequate air circulation around the enclosure

INADEQUATE AIR CIRCULATION

Air circulation around the Vess R2600 enclosure might be a more complex problem. Use the thermometer icons to help you locate the specific hot spot. Check for these conditions:

- Accumulated dust or objects blocking the fans
- Less than a minimum of 13 cm (5 inches) space between the back of the enclosure and the wall or other object
- Ambient temperature above 35°C (95°F) where the subsystem is operating
- Failed Cooling Unit

To cool down an enclosure:

- Correct any problems identified above.
- Power it down and let it sit for an hour or longer.

See “Shutting Down the Subsystem” on page 58.

If a Cooling Unit must be replaced, do not remove the failed unit until a replacement unit is available. A hole on the backplate created by a missing Cooling Unit affects air circulation which will cause controller units to overheat and shut down. If a replacement Cooling Unit is not available, leave the failed unit in place until one is available. See "Replacing a Cooling Unit" on page 178.



Important

In the event of a Cooling Unit failure, **DO NOT** remove the failed unit until there is a replacement available and on hand. A single functioning Cooling Unit is adequate for cooling the system as long as the failed Cooling Unit remains in place. Removing the failed unit without replacing it will adversely affect airflow within the enclosure resulting in critical overheating and shutdown of the enclosure.

POWER SUPPLIES

Vess R2600 subsystems are equipped with redundant power supplies. The advantage of N+1 power supplies is that should one fail, the other continues to power the subsystem until the faulty one can be replaced. The subsystem is capable of operation on N power supplies.

The power supplies are hot-swappable, meaning you can leave the subsystem running when you replace the bad one. Be careful, however, to remove the faulty power supply and not the good one, or the subsystem comes to an immediate stop and your data is unavailable until the subsystem is powered and booted again.

POWER FAN FAILURE

In the Vess R2600 subsystems, the power supply fans are used for cooling the power supply. When a power supply fan fails, you must replace the power supply.

See "Replacing a Power Supply" on page 177 for more information.

BATTERIES

The Cooling Units in the Vess R2600 subsystem use a battery for backup power to protect data in the cache. Should a power failure occur, the battery enables the cache to hold data up to 72 hours. The battery recharges during normal subsystem operation.

In most cases, installing a replacement battery corrects a marginal or failed condition. The battery is located inside the Cooling Unit housing. To replace a battery, first replace the Cooling Unit with the battery that needs to be replaced. See "Replacing a Cache Backup Battery" on page 179 Also see "Reconditioning a Battery" on page 62.

RAID CONTROLLER PROBLEMS

RAID controller problems include:

- "Maintenance Mode" on page 211
- "Finding and Correcting the Cause of the Problem" on page 211
- "Taking a RAID Controller out of Maintenance Mode" on page 211
- "Unsaved Data in the Controller Cache" on page 212

Controller problems occur when one of the controllers goes into maintenance mode.

MAINTENANCE MODE

For Vess R2600s with two RAID controllers, one of them enters **maintenance mode** in the event of:

- A difference of some kind between the two controllers (described below)
- An internal controller failure

When a controller enters maintenance mode, it goes offline and it displays N/A (not accessible) under Readiness Status.

You must find and correct the cause of the problem and then take the controller out of maintenance mode.

FINDING AND CORRECTING THE CAUSE OF THE PROBLEM

EXTERNAL CHECKS

Make the following external checks to your Vess R2600 subsystem. Be sure that:

- Both RAID controllers are present, fully inserted into their slots, and locked into place.
- The RAID controllers match, meaning both are Fibre Channel or both are iSCSI.
- All SAS expansion cables from the RAID controllers to external JBOD units in good condition and are securely connected.



Important

A disconnected SAS expansion cable causes the two RAID controllers to see a different set of configured drives. This condition is the most common cause of a controller entering maintenance mode.

INTERNAL CHECKS

If all external checks are OK, take the following actions:

1. Shut down the Vess R2600.
2. Remove one of the RAID controllers.

See “Replacing a RAID Controller – Dual Controllers” on page 180.

3. Restart the Vess R2600.
4. After the Vess R2600 is fully booted, view the controller information.

See page 60 (WebPAM PROe) or page 124 (CLU).

5. Observe and record the following information about the first controller:
 - SDRAM memory size
 - Hardware version
 - Firmware version
6. Shut down the Vess R2600.
7. Remove the first controller and install the second controller.
8. Repeat steps 3 through 6. Then compare your records.
9. Correct any differences between the two controllers.

See "Updating the Subsystem Firmware" on page 172.

TAKING A RAID CONTROLLER OUT OF MAINTENANCE MODE

If you shut down the Vess R2600 subsystem in the process of correcting the maintenance mode problem, the affected RAID controller boots into **normal mode** when the Vess R2600 restarts. No further action is required.

If you corrected the problem without shutting down the Vess R2600 subsystem, choose one of the following methods to take the controller out of maintenance mode:

- Restart the Vess R2600 subsystem
- Establish a serial connection, then use the CLI
- Establish a Telnet connection, then use the CLI

SERIAL CONNECTION

To clear maintenance mode using a serial connection:

1. Change your terminal emulation program settings to match the following specifications:
 - Bits per second: 115200
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: none

2. Start your PC’s terminal VT100 or ANSI emulation program.
3. Press **Enter** once to launch the CLI.

The login screen appears.

The following steps show the default Administrator user name and password. Use your own user name and password if you have changed these.

4. At the Login prompt, type **administrator** and press **Enter**.
5. At the Password prompt, type **password** and press **Enter**.

The CLI screen appears.

The prompt should display MAINTENANCE MODE@cli>. If the prompt displays your login name, such as administrator@cli>, log into the other controller.

6. At the MAINTENANCE MODE@cli> prompt, type **maintenance -a exit** and press **Enter**.

The controller reboots. The login screen again appears.

7. Close the Serial connection.

TELNET CONNECTION

This procedure requires you to know the IP address of the controller.

To clear maintenance mode using a Telnet connection:

1. Go to the command line prompt (Windows) or click the terminal icon (Linux), then run:

telnet 10.0.0.1 2300

The IP address above is only an example. 2300 is the Telnet port for Vess R2600.

The login screen appears.

The following steps show the default Administrator user name and password. Use your own user name and password if you have changed these.

2. At the Login prompt, type **administrator** and press **Enter**.
3. At the Password prompt, type **password** and press **Enter**.


The CLI screen appears.

The prompt should display MAINTENANCE MODE@cli>. If the prompt displays your login name, such as administrator@cli>, log into the other controller.

4. At the MAINTENANCE MODE@cli> prompt, type **maintenance -a exit** and press **Enter**.

The controller reboots. The Telnet session ends.

UNSAVED DATA IN THE CONTROLLER CACHE

The dirty cache LED (marked with the  icon) informs you that there is data in the cache that has not been saved to non-volatile memory. Such data is sometimes called “dirty,” not to suggest it is corrupted in some way but because it has not been saved to a physical drive.



Caution

If there is unsaved data in the controller’s cache, the dirty cache LED shines amber. During this time, do NOT power down the Vess R2600. Wait until the LED goes dark.

PHYSICAL DRIVE PROBLEMS

Physical drives are the foundation of data storage. A physical drive problem can affect your entire RAID system.

When a yellow !  icon or a red X  icon appears beside a physical drive, check the drive’s operational status:

1. Click the Device tab.
2. Click the Physical Drive icon.
3. Click the physical drive you want, then click the View button.

Look under Operational Status for the condition of the physical drive.

- **Offline** – Check the drive for:
 - PFA Condition** – Caused by a bad block or sector. See Note 1 below.
 - Stale Condition** – Caused by obsolete array information on the physical drive. See Note 2 below.
- **Not Usable** – This condition occurs when you have:
 - Two controllers in your RAID subsystem and a SATA drive without a SAS-to-SATA adapter. See Note 3 below.
 - A missing or defective SAS cable between the RAID subsystem and a JBOD expansion unit.
- **Drive Failed or Dead** – The physical drive cannot be repaired. You must replace the failed drive. See Note 4 below.

Note 1: Clear the error condition. Then the physical drive is available. See “Clearing a Stale or a PFA Condition”

Note 2: Identify the disk array to which the physical drive belongs. Then delete the disk array. If the error condition remains on the physical drive, clear the error condition.

Note 3: Obtain SAS-to-SATA adapters though PROMISE Technology, at <http://www.promise.com>.

Note 4: You can set the number of bad blocks tolerated before the controller marks a physical drive as Dead.

DISK ARRAY AND LOGICAL DRIVE PROBLEMS

Disk array and logical drive problems include:

- "Disk Array Degraded / Logical Drive Critical" on page 213
- "Disk Array Offline / Logical Drive Offline" on page 214
- "Repairing an Offline Disk Array or Logical Drive" on page 214
- "Rebuilding a Disk Array" on page 215
- "Incomplete Array" on page 215

Disk array problems typically result from a physical drive failure. The most common problem is a degraded disk array. The RAID controller can rebuild a degraded disk array.

DISK ARRAY DEGRADED / LOGICAL DRIVE CRITICAL

Disk arrays are made up of physical drives. Logical drives are created on the disk array.

When one of the physical drives in a disk array fails:



- The operational status of the **disk array** becomes Critical.
- The operational status of the **logical drives** becomes Critical or Degraded.
- The operational status of the **physical drive** becomes **Dead** or **Offline**.

WebPAM PROe reports these conditions in the following places:

- **Dashboard tab**
 - A yellow !  icon beside the disk arrays, logical drives, and physical drives under System Status.

Major event for the logical drive under Event Information.

Warning event for the physical drive under Event Information.

- **Device tab**
 - Front View** – Physical drives are shown **Dead** or **Offline** and marked with a red X  icon, or **Missing**.
 - Physical Drive View** – Physical drives are shown **Dead** or **Offline** and marked with a red X  icon, or **Missing**.

- Storage tab

Disk Array and Logical Drive are marked Critical with a yellow ! ⚠ icon.

RAID 6 and 60 logical drives are marked:

- Degraded with a yellow ! ⚠ icon when ONE physical drive is offline.
- Critical with a yellow ! ⚠ icon when TWO physical drives are offline.

RAID 0 logical drives show Offline status and a red X ❌ icon.

If there is no spare drive or unconfigured drive in the RAID system, you must provide the replacement drive. See “Installing Physical Drives” on page 16.

- Administration tab

Depending on your settings and availability of a replacement drive, your system automatically rebuilds the degraded disk array. See “Rebuilding a Disk Array” on page 215.

The system sends an Email message about the incident to subscribing users, depending on user settings. See “Setting User Event Subscriptions”.

DISK ARRAY OFFLINE / LOGICAL DRIVE OFFLINE

Disk arrays are made up of physical drives. Logical drives are created on the disk array. When a disk array and its logical drives go Offline, the data stored in the logical drives is no longer accessible.

RAID 0 logical drives go Offline when ONE physical drive is removed or fails.

RAID 1, 1E, 3, 5, 10, 30, and 50 logical drives go Offline when TWO physical drives are removed or fail.

RAID 6 and 60 logical drives go Offline when THREE physical drives are removed or fail.

WebPAM PROe reports these conditions in the following places:

- Dashboard tab

A red X ❌ icon appears beside the disk arrays, logical drives, and physical drives under System Status.

Major event for the logical drive under Event Information

Warning event for the physical drive under Event Information.

- Device tab

On Front View and Physical Drive View, physical drives are shown Dead, Offline, or Missing.

- Storage tab

Disk array and logical drives are marked with a red X ❌ icon.

- Administration tab

Under Background Activities, no Rebuild takes place. See Repairing, below.

The system sends an Email message about the incident to subscribing users, depending on user settings.

REPAIRING AN OFFLINE DISK ARRAY OR LOGICAL DRIVE

RAID 1, 1E, 3, 5, 6, 10, 30, 50, AND 60 LOGICAL DRIVES

If a fault-tolerant logical drive, RAID 1, 1E, 3, 5, 6, 10, 30, 50, and 60, goes Offline, it may be possible to recover your data.



Warning

Take no further corrective action until you have consulted with Technical Support!

RAID 0 LOGICAL DRIVES

If a logical drive based on a non-fault-tolerant disk array, RAID 0, goes offline, all of the data on the logical drive is lost.

To recreate your logical drive:

- Identify the failed physical drive.

See “Locating a Physical Drive”.
- Replace the failed drive.
- See “Installing Physical Drives” on page 16.
- If the disk array had more than one physical drive, delete the disk array and re-create it.

See “Deleting a Disk Array” and “Creating a Disk Array Manually”.

- Restore the data from your backup source.

REBUILDING A DISK ARRAY

When you rebuild a disk array, you are actually rebuilding the data on one physical drive.

- When a physical drive in a disk array fails and a spare drive of adequate capacity is available, the disk array begins to rebuild automatically using the spare drive.
- If there is no spare drive of adequate capacity, but the Auto Rebuild function is ENABLED, the disk array begins to rebuild automatically as soon as you remove the failed physical drive and install an unconfigured physical drive in the same slot. See “Making Rebuild Settings”
- If there is no spare drive of adequate capacity and the Auto Rebuild function is DISABLED, you must replace the failed drive with an unconfigured physical drive, then perform a Manual Rebuild. See “Rebuilding a Disk Array”



Important

If your replacement disk drive was formerly part of a different disk array or logical drive, you must clear the configuration data on the replacement drive before you use it. See “Clearing a Stale or a PFA Condition”

-

INCOMPLETE ARRAY

A more serious, but far less common problem is an **Incomplete Array**. An incomplete array results from a physical drive that fails or becomes missing during:

- RAID level migration
- Disk array transport

MIGRATION

Normally, if a physical drive or the controller fails during migration, the disk array goes critical, and you can rebuild it.

TRANSPORT

Transport is the action of moving the physical drives of a disk array:

- To different slots in the same enclosure
- From one enclosure to another

If a physical drive fails during a transport, or you do not move all of the physical drives to their new locations,

WebPAM PROe displays an incomplete array. When WebPAM PROe discovers an incomplete array, it displays a dialog box asking you to:

- Click the **OK** button to accept the incomplete array.
- Click the **Cancel** button to reject the incomplete array.

ACCEPTING AN INCOMPLETE ARRAY

Before you accept the incomplete array, be sure all of the physical drives are present and that their drive carriers are properly installed into the enclosure. See “Installing Physical Drives”

If you choose to accept the incomplete array:

1. Click OK in the incomplete array dialog box.
2. Check the operational status of the logical drives in the array.
 - If the logical drives are **Critical**, proceed with a rebuild.
 - If the logical drives are **Offline**, contact Technical Support.

3. Restore your data from a backup source.

If you choose NOT to accept the incomplete array:

1. Click **Cancel** in the incomplete array dialog box.
2. Do one of the following:
 - Delete the array. This action deletes all logical drives on the array.
 - Replace the missing physical drive.

CONNECTION PROBLEMS

Connection problems include:

- "Serial Connections" on page 216
- "Network Connections" on page 216
- "SAS Connections" on page 217
- "Browser Does Not Connect to WebPAM PROe" on page 217

Connection problems cause a majority of failures in almost any electrical system. While the installation of the cables and components was correct, they don’t function properly, or at all, because:

- A connector is dirty or corroded
- A connector is loose or damaged
- A cable looks OK outside but has an open circuit inside
- The wrong cable was used

Vess R2600s ship with a full set of new cables, as required for each specific model. Be sure to use these components because: 1.) They are the proper ones for your RAID subsystem, 2.) They are in brand-new condition, and 3.) You paid for them with the purchase of your subsystem.

SERIAL CONNECTIONS

Vess R2600 uses a serial connection for the command line interface (CLI) and the command line utility (CLU). After you set the IP address, you can access the CLI and CLU through a network connection, also. Normally, users prefer WebPAM PROe because of its graphic user interface. But the CLI and CLU can do the same jobs. And they work when your network connection is down.

For Vess R2600, you must use the CLI or CLU to set the Management Port IP address in order for WebPAM PROe to connect with it. See “Setting-up the Serial Connection” on page 35. This issue is discussed further under Network Connections, below. See “Making Serial Cable Connections” on page 29for more information on making the connection.

The CLI and CLU control and manage but they do not move data. They communicates through a RJ-11-to-DB9 serial data cable, supplied with the Vess R2600. You may choose not use the CLI or CLU often and want to disconnect and store the cable. Consider leaving it connected, so you know where it is the next time you need it.

NETWORK CONNECTIONS

Each RAID controller has an Ethernet (RJ45) management port connector on the back of the enclosure. This is a Gigabit Ethernet connector designed to connect to your network. The Vess R2600 becomes a node on your network like any other PC, server or other component with an IP address.

Vess R2600 ships from the factory IP addresses of 10.0.0.1, 10.0.0.2, and 10.0.0.3. You must change these addresses to ones that work on your network. You make the initial IP address setting using the CLI or CLU. See “Setting-up the Serial Connection” on page 44.

Management Port LEDs		
State	Activity	Connectivity
<i>Dark</i>	No activity	10BaseT
<i>Steady green</i>	—	100BaseT
<i>Flashing green</i>	Activity	—
<i>Amber</i>	—	1000BaseT

Note that Vess R2600’s virtual and maintenance ports can accept IP address assignments from a DHCP server. Use Vess R2600’s Command Line Utility (CLU) to enable this feature.

If you manually assigned an IP address to the Vess R2600 but there is a DHCP server on your network, there is a chance that the server might assign the Vess R2600’s IP address to another node. You might see a warning to this effect on your PC’s monitor. If this happens, WebPAM PROe may not be able to connect. See your network administrator to work out a suitable arrangement.

FIBRE CHANNEL CONNECTIONS

When there is a connection failure, use WebPAM PROe to verify that Vess R2600 sees the initiators. See “Viewing a List of FC Initiators on the Fabric” .

If Vess R2600 sees some initiators but not the one you want, the problem is most likely elsewhere in the loop or fabric. If Vess R2600 does not see any initiators:

- Check all of the Fibre Channel connections
- Verify that all nodes are properly connected and powered
- Verify that the fabric router or switch is properly connected powered

SAS CONNECTIONS

Faulty SAS expansion connections are suspected when the link port counter reports a large number of bad link errors.

Link errors can be caused by:

- Debris blocking the SAS cable connector
- A faulty SAS cable
- A faulty controller or I/O module SAS connector

BLOCKED CABLE CONNECTORS

To check for debris blocking the SAS cable connector:

1. Power down the RAID subsystem and JBOD units.
2. Remove the SAS cable and check all SAS connectors for debris.
3. Clean the connectors as required and reconnect the SAS cable.
4. Power up the subsystems and monitor the link port counter for changes in the rate of link error accumulation.

FAULTY CABLE

To check for a faulty SAS cable:

1. Power down the RAID subsystem and JBOD units.
2. Replace the SAS cable with a new one.
3. Power up the subsystems and monitor the link port counter for changes in the rate of link error accumulation.

FAULTY CONTROLLER OR I/O MODULE CONNECTOR

To check for a bad controller or I/O module SAS connector:

1. With the subsystems online and I/Os running, access the CLI via serial or Telnet.
2. At the command prompt, type the following command and press Enter.
administrator@cli> sasdiag -a errorlog -l expander -e 1 -i 1
3. At the command prompt, type the following command and press Enter.
administrator@cli> sasdiag -a errorlog -l c2cport

By interpreting the two error logs, you can verify which controller or I/O module SAS port is accumulating link errors.

BROWSER DOES NOT CONNECT TO WEBPAM PROE

If you successfully setup and connected to WebPAM PROe, then suddenly you can no longer connect, it might be the result of the following three conditions:

- DHCP is enabled on your Vess R2600’s virtual management port
- The DHCP server does not have a dedicated IP address for the Vess R2600
- The Vess R2600 restarted and your DHCP server assigned a new IP address

You must obtain the new IP Address for the virtual management port in order to direct your browser to the Vess R2600 and start WebPAM PROe.

To access the new IP address:

1. Start your PC’s terminal VT100 or ANSI emulation program.
2. Press Enter once to launch the CLI.
3. At the Login prompt, type administrator and press Enter.
4. At the Password prompt, type password and press Enter.
5. Type net and press Enter.
administrator@cli> net

=====						
CI	Port	Type	IP	Mask	Gateway	Link
=====						
Virtual		Mgmt	10.0.0.1	255.0.0.0	0.0.0.0	Up

The new virtual management port IP address and other network settings display.

6. Enter the new IP address into your browser to log into WebPAM PROe.

For more information, see “Making Serial Cable Connections” and “Logging into WebPAM PROe”

POWER CYCLING THE SUBSYSTEM

To power cycle a RAID subsystem means to:

- Shut down
- Turn off the power
- Turn on the power
- Restart

Power cycling is sometimes required as a remedial action but only when prompted by a message from software or when directed by Technical Support.

To power cycle the RAID subsystem:

1. Shut down the subsystem.

When the controllers shut down, your network connection is lost.

2. Manually turn off the system power by pressing the power button of the RAID unit for over five seconds.
3. Wait at least 10 seconds.
4. Manually press the power button of the JBOD units.
5. Manually press the power button of the RAID subsystem.
6. Wait no less than two minutes.
7. Do one of the following actions:
 - Open your browser and log into WebPAM PROe.
 - Re-establish your Telnet or SSH connection to the subsystem and open the CLU.

If you cannot log in immediately, wait 30 seconds and try again.



Important

If your RAID subsystem manages JBOD expansion units, always power on the JBOD expansion units first. Then power on the RAID subsystem.

EVENT NOTIFICATION RESPONSE

When you choose Event Notification, WebPAM PROe sends popup and/or email messages regarding its status.

The messages you see depend on your notification selection and what is currently happening in the Vess R2600.

The table below cites:

- **Reported Events** – Events that require you to take action
- **Corrective Actions** – The action you should take in response to the event

A list of event categories is shown below.

- | | |
|-----------------------------|----------------------------------|
| • Battery | • Physical Disk (Physical Drive) |
| • BBU | • PSU (Power Supply Units) |
| • Blade Server | • PSU Fans |
| • Cache | • RAID Level Migration |
| • Controller | • Rebuild |
| • CRC | • Redundancy Check |
| • Disk Array | • Resource |
| • Drive Interface | • SCSI |
| • Enclosure | • SEP |
| • Event Log | • Spare Check |
| • Fibre Channel | • Spare Drives |
| • Firmware Update | • SMART |
| • Host Interface | • Stripe Level Migration |
| • Initiator | • Synchronization |
| • JBOD | • Subsystem |
| • Logical Drive | • Transition |
| • Media Patrol | • Unknown |
| • Online Capacity Expansion | • Zoning |
| • Parity | |
| • PDM | |

Reported Event	Corrective Action
----------------	-------------------

Battery	
Battery is inserted	No action is required.
Battery charging has failed	Replace the battery.
Battery reconditioning has started	No action is required.
Battery reconditioning has been terminated	Replace the battery.
The write policy of writeback logical drive switched from writeback to writethru	Check the event log to see whether battery is re-conditioning.
The write policy of writeback logical drive switched from writethru to writeback	No action is required.
Battery is charging in high temperature	Monitor the condition. Contact Tech Support if the problem persists.
Battery cannot function with the enclosure or with the attached battery board	Wrong battery installed. Contact Tech Support for assistance.
Logical drive writeback cache maybe enabled without battery support	No action required.
Battery is fully charged	
Battery is not present	Install a battery or verify that the battery is properly connected.
Battery is not accessible	Connect the battery properly or replace the battery.
BBU	
BBU flushing has started	No action is required.
BBU flushing has ended	
BBU flushing has failed	Contact Tech Support if the condition persists.
Blade Server	
Blade Server Inserted	No action is required.
Blade Server Removed	
Cache	
Not available	Contact Tech Support.
Controller	
The controller parameter(s) are changed by user	No action is required.
The controller is reset by Watch Dog timer	Result of a firmware update. If the condition persists, replace the controller.
The controller has new crash information	Contact Tech Support.
The controller's heart beat has started	No action is required.
The controller's heart beat has stopped	
The partner controller's heart beat has started	
The partner controller's heart beat has stopped	
The partner controller's heart beat has skipped	
The controller's main scheduler has frozen	Contact Tech Support if the condition persists.

Controller has entered maintenance mode since configured physical disk seen by partner controller is not seen here	Verify that all SATA drives have an SAS-to-SATA adapter installed.
Controller has entered maintenance mode due to mismatch of physical disks types	Check and correct SAS cabling and connections as needed.
Controller has entered maintenance mode due to mismatch of physical disk WWN	Update to the latest firmware. If the condition persists, replace the controller.
Controller has entered maintenance mode due to mismatch of SATA Disks	Check and correct data cabling and connections as needed.
Controller has entered maintenance mode due to mismatch of Disk IDs	
Controller has entered maintenance mode since no physical disks are seen as seen by Partner controller	
Controller is started	No action is required.
Controller is set to Active Mode	
Controller is set to Standby Mode	
Controller Failed Over as partner is removed	Verify that the partner controller is properly installed and all cables are connected.
Controller Failed Over as heart beat stopped	
Controller Firmware mismatch with that of the partner controller	Auto Firmware synchronization upgrades or downgrades the firmware.
Controller set to Maintenance Mode because of hardware mismatch with partner (controller)	Compare controller types and amount of memory installed. Correct or update as needed.
Controller set to Maintenance Mode because of firmware mismatch with partner controller	Update this controller to the same firmware version as the partner controller.
Controller set to Maintenance Mode because Firmware is flashing in the partner controller	Exit out of Maintenance mode after firmware flashing is complete.
Controller set to Maintenance Mode because of flash image version mismatch with partner (controller)	Update this controller to the same flash image version as the partner controller.
Controller has been set to Maintenance mode because there is a mismatch in the Controller Model or Hardware version with that of the partner controller	Replace this controller with the same Model and Hardware version as the partner controller.
Controller has been set to Maintenance mode because there is a mismatch in the memory size with that of the partner controller	Replace this controller’s memory with the same memory size as the partner controller
Partner Controller has entered maintenance mode to protect user data since one of the configured physical drives was disconnected in the partner controller	Check and correct cable connections to external JBOD enclosures. Rebuild any critical logical drives. Back up array data. Replace the physical drive. Bring controller out of maintenance mode.
Controller was placed on reset during Fail Over processing	No action is required.
Partner Controller was placed on reset during Fail Over processing	

<i>Controller was reset as it was not able to join the running partner controller</i>	Verify that the controller is running. If the condition persists, replace the controller.
<i>The controller has reset because it encountered a firmware problem</i>	If resets happen frequently, update to new firmware or replace the controller.
<i>Controller temperature is above the warning threshold</i>	Check airflow around the Vess R2600. Check blowers and fans.
<i>The controller temperature is above controller critical threshold</i>	
<i>Controller temperature is within the normal range</i>	No action is required.
CRC	
<i>CRC error is detected while receiving CMD information unit</i>	If this message appears repeatedly, contact Tech Support.
<i>CRC error is detected during Data Out phase</i>	
Disk Array	
<i>New disk array has been created</i>	No action is required.
<i>Disk array has been deleted</i>	
<i>Disk array has been added</i>	
<i>Disk array has been removed</i>	
<i>Disk array settings have been changed</i>	
<i>Disk array is transport ready</i>	Remove physical drives in disk array and insert them into a different subsystem. To cancel Transport Ready Status, remove and reinsert the drives in their original slots.
Drive Interface	
<i>Drive-interface controller is found</i>	No action is required.
<i>Drive-interface controller is NOT found</i>	Restart the Vess R2600. If this message appears repeatedly, contact Tech Support.
<i>Drive-interface diagnostics has passed</i>	No action is required.
<i>Drive-interface diagnostics has failed</i>	Restart the Vess R2600. If this message appears repeatedly, contact Tech Support.
<i>Drive-interface controller has generated a general parity error</i>	If this message appears repeatedly, contact Tech Support.
<i>Drive-interface controller has generated a data parity error</i>	
Enclosure	
<i>Enclosure temperature is above the threshold</i>	Check blowers and fans.
<i>Enclosure temperature is above the warning threshold</i>	Check airflow around the Vess R2600. Check blowers and fans.
<i>Enclosure temperature is above the critical threshold</i>	
<i>Enclosure temperature is within the normal range</i>	No action is required.
<i>Shut down PSUs due to enclosure or controller temperature over threshold</i>	Shut down the Vess R2600
Event Log	

Event logging is enabled	No action is required.
Event logging is disabled	
Event log buffer is cleared in RAM	
Event log buffer is cleared in NVRAM	
Event log buffer is cleared in MDD	
Fibre Channel	
Fibre Channel controller has detected bus reset	If this message appears repeatedly, contact Tech Support.
Fibre Channel controller has received a “LUN reset” command.	No action is required.
Fibre Channel controller has encountered a fatal error	Restart the Vess R2600. If this message appears repeatedly, contact Tech Support.
Fibre Channel link is up	No action is required.
Fibre Channel link is down	
Fibre Channel controller settings have changed	
Firmware Update	
Firmware update is started	No action is required.
Firmware update is complete	
Firmware update is fail	Try the update again. If this message repeats, contact Tech Support.
Back-end expander firmware upgrade is started	No action is required.
Back-end expander firmware upgrade is completed	
Back-end expander firmware upgrade failed	Try the update again. If this message repeats, contact Tech Support.
Front-end expander firmware upgrade is started	No action is required.
Front-end expander firmware upgrade is completed	
Front-end expander firmware upgrade failed	Try the update again. If this message repeats, contact Tech Support.
Host Interface	
Host interface controller has detected bus reset	If this message appears repeatedly, contact Tech Support.
Host interface controller has encountered an unrecoverable error	Restart the Vess R2600. If this message appears repeatedly, contact Tech Support.
Host interface controller has received an “abort task” command.	No action is required.
Host interface controller has received an “abort task set” command.	
Host interface controller has received a “clear ACA” command.	If this message appears repeatedly, contact Tech Support.
Host interface controller has received a “clear task set” command.	No action is required.
Host interface controller has received a “LUN reset” command.	

Host interface controller is informed that the initiator has detected an error	If this message appears repeatedly, contact Tech Support.
Host interface controller has received illegal secondary identification	
Host interface controller has received a message parity error	
Host interface controller has received a bus reboot	
Host interface link is up	No action is required.
Host interface link is down	Check connections.
Host interface controller has encountered an unknown error	If this message appears repeatedly, contact Tech Support.
Host interface controller has encountered a system error	
Host interface controller has encountered a fatal error	Restart the Vess R2600. If this message appears repeatedly, contact Tech Support.
Host interface controller settings have changed	No action is required.
Host interface controller has received a ‘WARM reset’ command	If this message appears repeatedly, contact Tech Support.
Host interface controller has received a “COLD reset” command	
Host Interface controller, MU handshake failed	
Host Interface controller, HMU has stopped	
Host Interface controller, FMU has unloaded	
Initiator	
Initiator sent message for detecting an error	If this message appears repeatedly, contact Tech Support.
JBOD	
JBOD system connected	No action is required.
JBOD system either is removed or malfunctioned	Check Expander firmware and SAS connections.
Logical Drive	
Logical drive initialization has started	No action is required.
Logical drive Initialization is in progress	
Logical drive initialization has completed	
Logical drive initialization has paused	Resume the initialization when ready.
Logical drive initialization has resumed	No action is required.
Logical drive initialization has stopped	If this action was not intentional, check the logical drive’s status.
Logical drive initialization marks the logical drive offline	Replace the failed physical drive. Delete and recreate the logical drive.
Logical drive initialization is aborted due to an internal error.	Reduce system load on the Vess R2600.

Logical drive initialization is queued	No action is required.
Quick logical drive initialization has started	
Quick logical drive initialization has completed	
Quick logical drive initialization has paused	Resume the initialization when ready.
Quick logical drive initialization has resumed	No action is required.
Quick logical drive initialization has stopped	If this action was not intentional, check the logical drive's status.
Quick logical drive initialization marks the logical drive offline	Replace the failed physical drive. Delete and recreate the logical drive.
Quick logical drive Initialization is aborted due to an internal error	Reduce system load on the Vess R2600.
Quick logical drive initialization is queued	No action is required.
A new logical drive has been created	
Logical drive has been deleted	
Logical drive has been placed online	Check the state of the physical drives, replace any bad drives. Rebuild logical drive.
Logical drive has been placed online. Possible data loss	
Logical drive has been set to critical.	
Logical drive has been set to degrade	No action is required.
Rebuild marks the logical drive synchronized upon rebuild completion	
Logical drive settings has been changed through a user command	
One of the error tables of a logical drive has been cleared by the user	
Logical drive axle has been placed online	
Media Patrol	
Media patrol is started	No action is required.
Media patrol is in progress	
Media patrol is completed	
Media patrol is paused	Resume Media Patrol when ready.
Media patrol is resumed	No action is required.
Media patrol is stopped	If this action was not intentional, check the logical drive's status.
Media patrol is aborted due to an internal error.	Reduce system load on the Vess R2600.
Media patrol is queued	No action is required.
Media patrol is stopped internally	
Online Capacity Expansion	
Online capacity expansion has started	No action is required.
Online capacity expansion has completed	
Online capacity expansion has paused	Resume OCE when ready.
Online capacity expansion has resumed	No action is required.

Online capacity expansion has stopped	If this action was not intentional, check the logical drive’s status.
Online capacity expansion has encountered a physical disk error	Check the physical drive check table after OCE is finished.
Online capacity expansion is aborted due to an internal error.	Reduce system load on the Vess R2600.
Online capacity expansion is queued	No action is required.
Parity	
Parity error is detected during Data Out phase	If this message appears repeatedly, contact Tech Support.
PDM	
PDM is started	No action is required.
PDM is in progress	
PDM is completed	
PDM is paused	Resume PDM when ready.
PDM is resumed	No action is required.
PDM is stopped	If this action was not intentional, check the disk array’s status.
PDM is switched to rebuild.	Replace the dead physical drive or reinstall the missing drive.
PDM is stopped internally	The destination drive was removed or used for a rebuild.
Physical Disk (Physical Drive)	
Physical disk is marked online	No action is required.
Physical disk is marked offline	Replace the physical drive.
Physical disk is marked as DEAD.	
Physical disk has been reset	
Physical disk assigned as global spare	No action is required.
Global Spare has been deleted	
Physical Disk is no longer assigned as a global spare	
Physical disk assigned as dedicated spare	
Dedicated Spare has been deleted	
Physical Disk is no longer assigned as a dedicated spare	Insert the physical drive back into the system.
Physical disk has been inserted	
Physical disk has been removed	
Command on physical disk has been re-tried	If this message appears repeatedly, replace the physical drive
Physical disk ECC error is detected	Replace the physical drive.
Physical disk CRC error is detected	
Bad sector is found on physical disk	If this message appears repeatedly, replace the physical drive.
Error is detected in remap sectors	
Command times out on physical drive	
Physical disk negotiation speed is decreased.	
Previously configured disk is no longer found	Insert the physical drive back into the system.

A physical disk has encountered an unknown (non-ECC) media error.	If this message appears repeatedly, replace the physical drive.
A physical disk has encountered PFA condition	Clear the PFA condition. If this message appears repeatedly, replace the physical drive.
A configured dead physical drive has been inserted	Replace the physical drive.
A physical drive page 0 settings have been changed	No action is required.
A physical drive page 1 settings have been changed (SATA drives)	
A physical drive page 3 settings have been changed (SAS drives)	
Physical disk is marked as DEAD due to removal	Replace the physical drive.
Physical disk is marked as DEAD due to failure of reassign sectors command	
Physical disk is marked as DEAD due to PFA condition	
Physical disk is marked as DEAD due to forced offline state	
Physical disk seen by partner controller not seen here	Check and correct SAS connections. Verify that SAS-to-SATA adapters are installed on all SATA drives.
Single ported physical disk seen by Partner controller not seen here	Install an SAS-to-SATA adapter on the SATA drive.
Physical disk reported not ready	Replace the physical drive.
PSU (Power Supply Units)	
PSU is not inserted	Reinstall the power supply unit.
PSU is off	Turn on the power supply or plug in the power cable.
PSU is on	No action is required.
PSU is installed and turned on	
PSU is functional and turned on	
PSU is installed and turned off	Turn on the power supply or plug in the power cable.
PSU is functional and turned off	
PSU is malfunctioning and turned on	Replace the power supply unit.
PSU is malfunctioning and turned off	
PSU has been removed	
PSU 12V/5V/3.3V power is out of the threshold range	No action is required.
PSU 12V/5V/3.3V power is within the normal range	
PSU is critical. This may cause instability of the system	Check the power to the PSU. Verify that the correct PSU is installed.
PSU Fans	
PSU fan or blower has turned on	No action is required.
PSU fan or blower has turned off	
PSU fan or blower speed is increased	
PSU fan or blower speed is decreased	Replace the power supply.
PSU fan or blower is malfunctioning	

PSU fan or blower is inserted	No action is required.
PSU fan or blower is functioning normally	
PSU fan or blower is NOT installed	Check fans or blowers.
PSU fan status is unknown.	Check for proper installation and turn on the power supply. If the condition persists, replace the power supply.
RAID Level Migration	
RAID level migration is started	No action is required.
RAID migration is in progress	
RAID level migration is completed	
RAID level migration is paused	Resume migration when ready.
RAID level migration is resumed	No action is required.
RAID level migration is stopped	If this action was not intentional, check the logical drive’s status.
RAID level migration has encountered a physical disk error	Check the disk drive check table after migration and replace disk drive as needed.
RAID level migration is aborted due to an internal error.	Reduce system load on the Vess R2600.
RAID level migration is queued	No action is required.
Migration has detected stale NV Watermark	Wait to see if the watermark clears.
Migration has cleared stale NV Watermark	No action is required.
Array was made incomplete due to missing NV Watermark	If the array is online, try migration again. If the array is offline, delete and recreate the array.
User has accepted Incomplete Array. (Caused by a missing NV Watermark)	Rebuild the disk array.
Rebuild	
Rebuild is started	No action is required.
Rebuild is in progress	
Rebuild is completed	
Rebuild is paused	Resume rebuild when ready.
Rebuild is resumed	No action is required.
Rebuild is stopped	If this action was not intentional, check the logical drive’s status.
Rebuild stopped internally	Contact Tech Support.
Rebuild is aborted	Reduce system load on the Vess R2600.
Rebuild is queued	No action is required.
Auto rebuild cannot start	Install a target physical drive of adequate capacity.
Redundancy Check	
Redundancy Check is started	No action is required.
Redundancy Check is completed	
Redundancy Check is paused	Resume Redundancy Check when ready.
Redundancy Check is resumed	No action is required.
Redundancy Check is stopped	
Redundancy Check is aborted due to internal error	Reduce system load on the Vess R2600.

Redundancy Check encountered inconsistent block(s)	Check the disk drive check table after RC and replace disk drive as needed.
Redundancy Check task is queued	No action is required.
Redundancy check is in progress	
Redundancy Check task is stopped internally	Restore the disk array to functional status.
Redundancy check is started on unsynchronized logical drive	No action is required.
Resource	
Resource is NOT available	Reduce system load on the Vess R2600.
SCSI	
SCSI host interface controller settings have changed	No action is required.
SEP	
SEP is found	No action is required.
SEP is NOT found	Insert or replace SEP hardware.
SEP I2C device access failure	If this message appears repeatedly, contact Tech Support.
SEP I2C device access recovered from failure	
Spare Check	
Spare check started on the given spare drive	No action is required.
Spare check completed successfully on the given spare drive	
Spare Drives	
Physical disk assigned as global spare	No action is required.
Physical disk is no longer assigned as global spare	
Global Spare has been deleted	
Physical disk assigned as dedicated spare	
Physical disk is no longer assigned as dedicated spare	
Dedicated Spare has been deleted	
SMART	
SMART error is received	If this message appears repeatedly, replace the physical drive.
Stripe Level Migration	
Stripe Level migration is started	No action is required.
Stripe Level migration is completed	
Stripe Level migration is paused	Resume SLM when ready.
Stripe Level migration is resumed	No action is required.
Stripe Level migration is stopped	If this action was not intentional, check the logical drive’s status.
Stripe Level migration has encountered a physical disk error	Check the physical drive check table after OCE is finished.
Stripe Level migration is aborted due to an internal error.	Reduce system load on the Vess R2600.
Stripe Level migration is queued	No action is required.
Synchronization	
Synchronization is started	No action is required.

<i>Synchronization is completed</i>	No action is required.
<i>Synchronization is paused</i>	Resume synchronization when ready.
<i>Synchronization is resumed</i>	No action is required.
<i>Synchronization is stopped</i>	
<i>Synchronization is aborted due to an internal error.</i>	Reduce system load on the Vess R2600.
<i>Synchronization is queued</i>	No action is required.
<i>Synchronization is stopped internally</i>	
<i>Subsystem</i> (Vess R2600)	
<i>The Subsystem is started</i>	No action is required.
<i>The Subsystem is stopped</i>	
<i>Subsystem parameter(s) are changed by user</i>	
<i>System is set to Redundant mode</i>	
<i>System is set to Critical mode</i>	Check controller operation.
<i>Enclosure has only N PSU left</i>	Check to make sure each enclosure has three or more operational PSU. (N=2 for Vess R2600 and Vess J2600 2000 models)
<i>Enclosure has only N-1 PSU left</i>	
<i>System is set to Non-Redundant mode</i>	If your system has two controllers, check controller operation.
<i>Transition</i>	
<i>Transition is started</i>	No action is required.
<i>Transition is completed</i>	
<i>Transition is paused</i>	Resume transition when ready.
<i>Transition is resumed</i>	No action is required.
<i>Transition is stopped</i>	If this action was not intentional, check the disk array's status.
<i>Transition was switched to rebuild</i>	Replace the dead physical drive or reinstall the missing drive.
<i>Unknown</i>	
<i>Unknown priority reason is detected</i>	If this message appears repeatedly, contact Tech Support.
<i>Zoning</i>	
<i>Zoning permission settings with the expander has been reset to defaults</i>	No action is required.
<i>Zoning expander has been rebooted.</i>	
<i>Zoning permission settings with the expander different than expected</i>	Settings have been updated correctly. No action is required.

CHAPTER 9: SUPPORT

This chapter contains the following topics:

- "Frequently Asked Questions" on page 225
- "Contacting Technical Support" on page 228
- "Returning the Product For Repair" on page 230

FREQUENTLY ASKED QUESTIONS

PHYSICAL DRIVES

What kind of disk drives can I use with Vess R2600?

Vess R2600 supports:

- 3.5-inch and 2.5-inch form factor
- Hard disk drives (HDDs) and solid state drives (SSDs)
- SAS, 6 Gb/s and 3 Gb/s
- SATA, 6 Gb/s and 3 Gb/s
- Supports any mix of SAS and SATA drives simultaneously in the same enclosure

For a list of compatible drives, go to PROMISE support:

<http://www.promise.com/support/>.

Vess R2600 does not support Parallel ATA (PATA) disk drives.

Why are all the disk drives in my JBOD marked Dead?

This condition happens when the JBOD expansion subsystem is disconnected from the RAID subsystem, powered off while the RAID subsystem is running, or powered on after the RAID subsystem was powered on. Use the force online function to restore the disk drives. See "Clearing a Stale or a PFA Condition" on page 130.

See "Making Webserver Settings" on page 79.

I can access the Vess R2600 over my company's intranet. But I can't access it from an outside Internet connection. How do I make the Internet connection work?

This condition is not related to Vess R2600, but is due to your firewall and network connection protocol. Contact your MIS Administrator.

Why can a RAID 1 logical drive on Vess R2600 consist of only two disk drives?

RAID 1 logical drives work in mirrored physical drive pairs. You could create up to eight RAID 1 logical drives. Or you can create a single RAID 1E or RAID 10 logical drive with data mirroring and up to 16 physical drives.

See "Installing Physical Drives" on page 16 and "RAID Levels" on page 184 for more information on the number of physical drives you can use for each RAID level.

Are logical drives on Vess R2600 limited to 2 TB?

No. But verify that your operating system supports logical drives over 2 TB.

Also, for the operating system to recognize the full capacity of logical drives over 2 TB, you must specify a sector size of 1 KB or larger when you create the logical drive. See "Sector Size" on page 192 for more information.

How can I be sure everything is working OK on the Vess R2600?

Locally: The Vess R2600 enclosure has LEDs on the front to monitor the status of power, field replaceable units (FRUs) and logical drives. When these are green, Vess R2600 is functioning normally. See "Front Panel LEDs" on page 202.

Remotely: Check the Dashboard tab in WebPAM PROe. See "WebPAM PROe Reports a Problem" on page 208.

If there are no yellow ! or red X warning icons displayed, Vess R2600 is functioning normally.

Can Vess R2600 run using just one power supply?

The Vess R2600 must use at least two power supplies.

Vess R2600 must use at least TWO power supplies. What happens if a fan fails?

If the system reports a fan malfunction, contact Technical Support (see PROMISE Technical Support provides several support options for PROMISE users to access information and updates. We encourage you to use one of our electronic services, which provide product information updates for the most efficient service and support.) immediately to schedule replacement of the suspect power supply as soon as possible. Running the unit in this condition for more than three weeks may shorten subsystem life and void your warranty.

What happens if a logical drive goes critical?

On the front of Vess R2600, the logical drive LED turns amber and the buzzer sounds (if enabled). See “Vess R2600 is Beeping” on page 201 and “LEDs Display Amber or Red” on page 202.

Vess R2600’s Netsend service does not report all events to Windows PCs.

This condition results from a shortcoming in Windows Messenger that causes miscommunication with Netsend. PROMISE is developing a workaround at the time of this writing. Note that all events are correctly reported in the Event Viewer.

STARTUP

How can I tell when the Vess R2600 has fully booted?

When the Vess R2600 is fully booted up, the Power and FRU LEDs light up green. If a disk array is present, the Logical Drive LED lights up green also. The Controller heartbeat LED blinks green once per second for five seconds, goes dark for ten seconds, then blinks green once per second for five seconds again. See “Front Panel LEDs” on page 202.

LOGGING IN

Why am I not able to log in?

Check the spelling and case. User names and passwords are case sensitive.

I have entered correct user name and password, but still I am not able to log in.

The Administrator may have deleted or disabled your user name.

The login screen says “Login failed: the requested service is busy.”

The subsystem might still be booting or rebooting. Dual controller subsystems take longer because the controllers boot individually, then they synch to each other. Wait a few moments, then try again.

CONNECTION

Why can’t I connect to my RAID System?

Be sure you are using the correct IP address and entry text for the Vess R2600 RAID subsystem. For more information, see “Logging into WebPAM PROe” on page 50.

I verified the IP address and entry text but I still cannot connect.

Check the physical network connections on the Vess R2600 RAID subsystem. If these are OK, report the problem to your network administrator.

I can access the Vess R2600 over my company’s intranet. But I can’t access it from an outside Internet connection. How do I make the Internet connection work?

This condition is not related to the Vess R2600 or WebPAM PROe. The problem is caused by your firewall or network connection protocol. Contact your network administrator for help.

I tried to log into WebPAM PROe but my browser showed the message “cannot be displayed.” What is the problem?

The browser decided prematurely that WebPAM PROe was not responding. Click the browser’s **Refresh** button. This action usually brings up the login screen.

TIMEOUTS

WebPAM PROe was working OK. But then it timed out. What do I do now?

WebPAM PROe times out when 24 minutes have passed with no user activity. User activity means any action you do in WebPAM PROe to view or manage the subsystem. This feature is included for security purposes.

Have your administrator change the timeout interval.

Or to prevent WebPAM PROe from timing out, periodically click the interface with your mouse.

EMAIL MESSAGES

Why don’t I receive email messages from WebPAM PROe?

Check your User Event Subscription and User Email settings. See page 71. If these are correct, see your network administrator for assistance with the mail server setup, email accounts, and other issues.

USER MANAGEMENT

Why can’t I create a new User?

Only the Administrator or a **Super User** can create a User.

If you are the Administrator or a **Super User** and cannot create a User, be sure the user name is not already in use.

If you still cannot create a User, contact Technical Support. See PROMISE Technical Support provides several support options for PROMISE users to access information and updates. We encourage you to use one of our electronic services, which provide product information updates for the most efficient service and support..

Can I change my access rights?

Only the Administrator or a **Super User** can change user access rights. See page 70.

LOCK

Person “xyz” set the lock and is not available. How do I unlock it?

Ask your Administrator to release the lock.

Note: The lock automatically releases after the set amount of time has passed.

CREATING A DISK ARRAY OR LOGICAL DRIVE

Why can't I see all RAID Levels in RAID Level drop-down menu?

The selection of RAID Levels shown depends on number of physical drives available to the controller. For example, if there are only two physical drives, then you cannot see RAID 10, which requires four physical drives or RAID 50 that requires at least six. See "RAID Levels" on page 184.

Why can't I create more than one logical drive on my disk array?

If your logical drive takes up the entire capacity of the disk array, there is no room for another logical drive. Backup your important data, then delete the existing logical drive and create multiple smaller logical drives on the disk array. See "Deleting a Logical Drive" on page 98 and 139 and "Creating a Logical Drive Manually" on page 97.

DISK ARRAY DEGRADED

According to WebPAM PROe, my disk array is degraded. What am I supposed to do?

If the Auto Rebuild option is enabled and a hot spare drive is available, your disk array begins rebuilding automatically. Also see "Disk Array Degraded / Logical Drive Critical" on page 213

When an disk array becomes degraded, can I still access the data on it?

Yes, but reads and writes take longer while rebuilding is in progress.

DELETING A DISK ARRAY OR LOGICAL DRIVE

Why can't I select Delete Disk Array or Delete Logical Drive?

You must have Power or **Super User** Rights to delete a disk array or logical drive. See your Administrator about upgrading your access rights, if necessary. Also see "Making User Settings" on page 70.

Can I delete a Logical Drive without deleting the Disk Array?

Yes. See "Deleting a Logical Drive" on page 98 and 139.

REBUILDING A DISK ARRAY

I replaced a failed physical drive with a used but known-good drive. The system does not rebuild to it. Why not?

The replacement drive was previously used in a different disk array or logical drive. You must clear (erase) the Reserve Sector of the replacement disk drive before the system can rebuild to it. "Clearing a Stale or a PFA Condition" on page 130.

MIGRATING A DISK ARRAY

When I try to migrate a disk array from one RAID level to another, why doesn't the controller let me do it?

Most RAID levels have a minimum and maximum number of physical drives. Be sure you have the correct number of drives available for the target disk array. See "RAID Levels" on page 184 and "RAID Level Migration" on page 189.

I want to add two more drives to my RAID 1 disk array. Why won't the controller let me migrate it?

A RAID 1 disk array uses only two disk drives. In this case, you can add two physical drives, then convert to a RAID 10 disk array. Most RAID levels have a minimum and maximum number of disk drives. See "RAID Levels" on page 184.

MEDIA PATROL AND PDM

Media Patrol and PDM are enabled on my system. But they never report anything.

This is a good sign. Media Patrol and PDM operate transparently until they find a problem on a physical drive.

CONTACTING TECHNICAL SUPPORT

PROMISE Technical Support provides several support options for PROMISE users to access information and updates. We encourage you to use one of our electronic services, which provide product information updates for the most efficient service and support.

PROMISE E-Support: <https://support.promise.com>

PROMISE web site: <http://www.promise.com//>

When you contact us, please have the following information available:

- Product model and serial number
- BIOS, firmware, and driver version numbers
- A description of the problem / situation
- System configuration information, including: motherboard and CPU type, hard drive models, SAS/SATA/ATA/ATAPI drives & devices, and other controllers.

UNITED STATES

580 Cottonwood Drive

Milpitas, Ca 95035, USA

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

AUSTRALIA

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

EMEA

NETHERLANDS

Science Park Eindhoven 5228

5692 EG Son, The Netherlands

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

AUSTRIA

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

FRANCE

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

GERMANY

Europaplatz 9

44269 Dortmund, Germany

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

SWEDEN

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

SWITZERLAND ITF

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com//>

NORWAY *ITF*

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

BELGUIM

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

LUXEMBOURG

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

UNITED KINGDOM

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

TAIWAN

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

CHINA

Room 1108, West Wing, Shi Chuang Plaza, 22 Information Road

Shangdi IT Park, Haidian District, Beijing 100085

Fax: 86-10-8857-8015

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

KOREA

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

HONG KONG

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

SINGAPORE

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

JAPAN

3F, Mura Matsu Bldg, 3-8-5, Hongo Bunkyo-ku

Tokyo 113-0033, Japan

Technical Support (E-Support): <https://support.promise.com>

Web site: <http://www.promise.com/>

LIMITED WARRANTY

PROMISE Technology, Inc. ("PROMISE") warrants that this product, from the time of the delivery of the product to the original end user:

- a) all components, except the cache backup battery, for a period of three (3) years;
- b) the cache backup battery, for a period of one (1) year;
- c) will conform to PROMISE’s specifications;
- d) will be free from defects in material and workmanship under normal use and service.

This warranty:

- a) applies only to products which are new and in cartons on the date of purchase;
- b) is not transferable;
- c) is valid only when accompanied by a copy of the original purchase invoice.
- d) Is not valid on spare parts.

This warranty shall not apply to defects resulting from:

- a) improper or inadequate maintenance, or unauthorized modification(s), performed by the end user;
- b) operation outside the environmental specifications for the product;
- c) accident, misuse, negligence, misapplication, abuse, natural or personal disaster, or maintenance by anyone other than a PROMISE or a PROMISE-authorized service center.

DISCLAIMER OF OTHER WARRANTIES

This warranty covers only parts and labor, and excludes coverage on software items as expressly set above. Except as expressly set forth above, PROMISE disclaims any warranties, expressed or implied, by statute or otherwise, regarding the product, including, without limitation, any warranties for fitness for any purpose, quality, merchantability, non-infringement, or otherwise. PROMISE makes no warranty or representation concerning the suitability of any product for use with any other item. You assume full responsibility for selecting products and for ensuring that the products selected are compatible and appropriate for use with other goods with which they will be used.

PROMISE does not warrant that any product is free from errors or that it will interface without problems with your computer system. It is your responsibility to back up or otherwise save important data before installing any product and continue to back up your important data regularly.

No other document, statement or representation may be relied on to vary the terms of this limited warranty. PROMISE’s sole responsibility with respect to any product is to do one of the following:

- a) replace the product with a conforming unit of the same or superior product;
- b) repair the product.

PROMISE shall not be liable for the cost of procuring substitute goods, services, lost profits, unrealized savings, equipment damage, costs of recovering, reprogramming, or reproducing of programs or data stored in or used with the products, or for any other general, special, consequential, indirect, incidental, or punitive damages, whether in contract, tort, or otherwise, notwithstanding the failure of the essential purpose of the foregoing remedy and regardless of whether PROMISE has been advised of the possibility of such damages. PROMISE is not an insurer. If you desire insurance against such damage, you must obtain insurance from another party.

Some states do not allow the exclusion or limitation of incidental or consequential damages for consumer products, so the above limitation may not apply to you.

This warranty gives specific legal rights, and you may also have other rights that vary from state to state. This limited warranty is governed by the State of California.

YOUR RESPONSIBILITIES

You are responsible for determining whether the product is appropriate for your use and will interface with your equipment without malfunction or damage. You are also responsible for backing up your data before installing any product and for regularly backing up your data after installing the product. PROMISE is not liable for any damage to equipment or data loss resulting from the use of any product.

RETURNING THE PRODUCT FOR REPAIR

If you suspect a product is not working properly, or if you have any questions about your product, contact our Technical Support staff, and be ready to provide the following information:

- Product model and serial number (required)
- Return shipping address
- Daytime phone number

- Description of the problem
- Copy of the original purchase invoice

The technician helps you determine whether the product requires repair. If the product needs repair, the technician issues an RMA (Return Merchandise Authorization) number.

Important

Obtain an RMA number from Technical Support **before** you return the product and write the RMA number on the label. The RMA number is essential for tracking your product and providing the proper service.

Return ONLY the specific product covered by the warranty. Do not ship cables, manuals, CDs, etc.

USA and
Canada:

PROMISE Technology, Inc.
Customer Service Dept.
Attn.: RMA # _____
47654 Kato Road
Fremont, CA 94538

Other
Countries:

Return the product to your dealer or retailer.
Contact them for instructions before shipping the product.

You must follow the packaging guidelines for returning products:

- Use the original shipping carton and packaging
- Include a summary of the product’s problem(s)
- Write an attention line on the box with the RMA number
- Include a copy of your proof of purchase

You are responsible for the cost of insurance and shipment of the product to PROMISE. Note that damage incurred due to improper transport or packaging is not covered under the Limited Warranty.

When repairing returned product(s), PROMISE may replace defective parts with new or reconditioned parts, or replace the entire unit with a new or reconditioned unit. In the event of a replacement, the replacement unit is under warranty for the remainder of the original warranty term from purchase date, or 30 days, whichever is longer.

PROMISE pays for standard return shipping charges only. You must pay for any additional shipping options, such as express shipping.

APPENDIX A: USEFUL INFORMATION

The appendix covers the following topics:

- SNMP MIB Files (below)
- Adding a Second RAID Controller
- Installing a Second RAID Controller

SNMP MIB FILES

PROMISE supplies two MIB files to integrate the Vess R2600 subsystem into your SNMP system. These files are in the SNMP folder on the Software CD.

The MIB files are:

- FCMGMT-MID.mib
- ISCSI-MIB.mib
- raidv4.mib

For help loading the MIB files, see the instructions that came with your MIB browser.

ADDING A SECOND RAID CONTROLLER

If your Vess R2600 subsystem shipped with one RAID controller, you can add a second RAID controller. The second controller must have:

- The same firmware version as the currently installed controller
- The same amount of SDRAM as the currently installed controller

To obtain information for the currently installed RAID controller:

1. Click the Device tab.
2. Click the **Component List** icon.
3. Click the Controller and click the View button.
4. On the Information tab, note the Firmware Version.
5. Click the Advanced information tab.
6. Note the Slot 1 and Slot 2 Memory Size.
7. Contact contact PROMISE Technical Support to order your second RAID controller.



PROMISE Technical Support prepares the new RAID controller with firmware and SDRAM to match the existing RAID controller in your Vess R2600 subsystem.

INSTALLING A SECOND RAID CONTROLLER

To install a second RAID controller in your Vess R2600 subsystem:

1. Shut down the subsystem.
2. Remove the blank cover from the right RAID controller slot.
3. Carefully slide the new RAID controller into the slot until the handle locks in place.
4. Attach your data and management cables to the new controller, as needed.

See "Making Management and Data Connections" on page 18 for cable connection information.

5. Power up the subsystem and launch WebPAM PROe.
6. In WebPAM PROe, click the Dashboard tab and look under System Status.
 - If the new controller has a green check
 -  icon, the installation is completed. Go to “New Settings for Dual Controllers”
 - If the new controller has a yellow !  icon, one of the RAID controllers went into **maintenance mode** because its firmware or memory do not match the other RAID controller. See “RAID Controller in Maintenance Mode,” below.

RAID CONTROLLER IN MAINTENANCE MODE

To manage a RAID controller in maintenance mode:

1. Click the Administration tab.
2. Click the Firmware Update icon.
3. Click the Controller Firmware Update option.
4. Compare the Firmware version on Controller 1 and Controller 2.
 - If the firmware versions are different, go to “Updating the Subsystem Firmware”
 - If the firmware versions match, contact PROMISE Technical Support for help installing the correct memory into the RAID controller.

NEW SETTINGS FOR DUAL CONTROLLERS

With the second controller successfully installed, make the following settings:

- Redundancy Type – Set to Active-Active or Active-Standby.
- See "Making Subsystem Settings (CLU)" on page 122.

- LUN Affinity – If you choose Active-Active redundancy.
See “Making Controller Settings” on page 86 or page 216.



Note

The Vess R2600 subsystem boots its RAID controllers sequentially. With a second controller installed, your subsystem takes about a minute longer to boot. This condition is normal.

DUAL CONTROLLERS AND SATA DRIVES

If your Vess R2600 subsystem has SATA disk drives installed, you must install a SAS-to-SATA adapter on each of the SATA drives.

Without the SAS-to-SATA adapter, SATA drives display a red X 🛑 icon and **Not Usable** status.

Obtain SAS-to-SATA adapters from PROMISE Technology at <http://www.promise.com>.

SAS drives do not require adapters.

Also see “Installing Physical Drives” on page 16 and "Contacting Technical Support" on page 228.

APPENDIX B: MULTIPATHING ON WINDOWS

The appendix covers the following topics:

- Before You Begin
- Installing PerfectPath
- Verifying Installation
- Running Perfect Path View
- Monitoring Your LUNs and Paths
- Features and Settings
- Troubleshooting
- Updating PerfectPath
- Repairing PerfectPath
- Removing PerfectPath

PerfectPath is a multipathing software designed for use with PROMISE Vess R2600 RAID subsystem products and includes:

- GUI – Graphic user interface, PerfectPath View, for easy monitoring and settings.
- DSM – Device-Specific Module driver.
- **Events Service** – Notification service posts events to the application log.

PerfectPath supports Fibre Channel and Serial Attached SCSI (SAS) technologies.

PerfectPath runs on Windows Server 2008 and 2008 R2 operating systems, on both x86 and x64 platforms.

For a list of supported OSes, download the latest compatibility list from PROMISE support: <http://www.promise.com/support/>.

BEFORE YOU BEGIN

Before you install PerfectPath on your Windows Host PC, you must:

- Install your Fibre Channel or SAS HBA cards and their device drivers.
- Close all computer and storage management applications, including Computer Management, Device Manager, Disk Management, and the Registry Editor.



Note

If you have a complex configuration, such as multiple HBAs connected with multiple LUNs and paths to your PC, installation can take a long time. You can choose to temporarily disconnect your storage, install PerfectPath, then reconnect your storage to reduce installation time.

INSTALLING PERFECTPATH

To install the PerfectPath software:

1. Download the PerfectPath installer file from PROMISE support:
<http://www.promise.com/support/> and save the installer file to your Windows desktop.
2. Double-click the PerfectPath.exe installer file to start the installer.
3. In the Welcome screen, click the Next button.
4. In the License Agreement screen, click the “I accept the terms of this license agreement” option, then click the Next button.
5. In the Close All Disk Management Applications screen, click the Next button.
6. In the Ready to Install the Program screen, click the Install button.
7. Optional. If the installer displays a Security Alert message about an unsigned driver, click the Yes button to continue installation.

The software files install onto the system drive in the Program Files\Promise\PerfectPath folder. There is no optional install location.

8. In the Install Completed screen, click the Finish button.
9. In the Restart message box, click the Yes button to restart your PC.



Important

Save the PerfectPath installer file in case you need to repair your PerfectPath software in the future. See “Repairing PerfectPath”

- 10.

VERIFYING INSTALLATION

Before you can verify PerfectPath installation:

- Your Host PC must have multiple data-path connections to the Vess R2600 subsystem.
- The Vess R2600 must be fully booted.
- The Vess R2600 must have at least one logical drive.

See "Making Management and Data Connections" on page 18 for information about making data connections.

See “Creating a Disk Array Manually” on page 92 or “Creating a Disk Array” on page 132 for information about creating RAID arrays and logical drives.

You can verify Perfect Path installation on the Host PC in one of three ways:

- Start Menu
- Services List
- Device Manager

START MENU

To verify PerfectPath installation in the Start menu:

From the Start menu, choose All Programs > PerfectPath > PerfectPath View.

The PerfectPath View software starts.

SERVICES LIST

To verify PerfectPath installation in the Services list:

1. From the Start menu, right-click the Computer icon and choose Manage from the popup menu.
2. In the Server Management tree, click the + icon beside Configuration.
3. Click the Services icon.
4. In the Services window, look for the PerfectPath Events Service.

If the PerfectPath Events Service is present, PerfectPath has been installed. The Service should be Started and set to Automatic on the Local System.

DEVICE MANAGER

To verify PerfectPath installation in the Device Manager:

1. From the Windows desktop, right-click the Computer icon and choose Manage from the drop-down menu.
2. In the Server Management tree, click the + icon beside Diagnostics.
3. Under Diagnostics, click the **Device Manager**.
4. In the Device Manager window, click Disk drives.
5. Under Disk drives, look for “Promise Vess R2600 Multi-Path Disk Device” in the Disk drives list.

RUNNING PERFECT PATH VIEW

Running PerfectPath View includes these functions:

- Starting PerfectPath View
- Quitting PerfectPath View

STARTING PERFECTPATH VIEW


To start PerfectPath View:

From the Start menu, choose All Programs > PerfectPath > PerfectPath View.

The PerfectPath View window opens.

QUITTING PERFECTPATH VIEW

To quit the PerfectPath View application, do one of the following actions:

- From the System menu, choose Exit.
- Click the Close  icon on the PerfectPath View window.

MONITORING YOUR LUNs AND PATHS

Monitoring your LUNs and Paths includes these functions:

- Viewing LUN Properties
- Viewing Path Properties
- Viewing LUN Performance Statistics
- Viewing Path Performance Statistics
- Viewing Events
- Clearing Path Statistics

VIEWING LUN PROPERTIES

To view a list of all LUNs:

1. Click a Server  in Tree View.
2. Click the Properties tab.

The Properties tab reports:

- **System** – Name, OS type, and version
- **LUNs** – Name, size, serial number, and load balance policy

Move the scroll bar or expand the window to see all of the reported information.

To view a single LUN and all of its Paths:

1. Click the LUN  in Tree View.
2. Click the Properties tab.

The Properties tab reports:

- **LUNs** – Name, size, and load balance policy
- **Paths** – Path ID, state, and adapter name


Move the scroll bar or expand the window to see all of the reported information.

See also:

- “Load Balance Policy”
- “Refreshing the Objects”

VIEWING PATH PROPERTIES

To view Path properties:

1. Click a Path  in Tree View.
2. Click the Properties tab.

The Properties tab reports:


- SCSI Port number
- SCSI Path ID
- SCSI Target
- SCSI LUN

See also:

- “Load Balance Policy”
- “Refreshing the Objects”

VIEWING LUN PERFORMANCE STATISTICS

To view performance statistics for a LUN:

1. Click the LUN  in Tree View.
2. Click the Performance tab.

The Performance tab reports the state and cumulative counts for each path to that LUN.

Move the scroll bar or expand the window to see all of the reported statistics.

The Performance tab reports the following data for each path:

Path ID	Bytes Written
State (Active or not)	Non-IO Requests
Read Requests	Queue Depth
Write Requests	Retries Count
Bytes Read	Failure Count

An Active state indicates this path is available to handle I/O requests.

If Active does not appear, the path is designated as Standby.


Active and Standby states are determined by Load Balance Policy.

See also:

- “Viewing Path Performance Statistics”
- “Load Balance Policy”

VIEWING PATH PERFORMANCE STATISTICS

To view performance statistics for a Path:

1. Click a Path  in Tree View.
2. Click the Performance tab.

The Performance tab reports the state and cumulative counts for a specific path.

Move the scroll bar or expand the window to see all of the reported statistics.

The Performance tab reports the following data for each path:

Path ID	Bytes Written
State (Active or not)	Non-IO Requests
Read Requests	Queue Depth
Write Requests	Retries Count
Bytes Read	Failure Count

An Active state indicates this path is available to handle I/O requests.

If Active does not appear, the path is designated as Standby.

Active and Standby states are determined by Load Balance Policy.

See also:

- Viewing Path Performance Statistics
- “Load Balance Policy”

VIEWING EVENTS

Click the Events tab to view MPIO related events. The data includes:

Type – Error, Warning, or Information	Server
Time	Message
Date	

Move the scroll bar or expand the window to see all of the reported information.

Use this information to verify that settings changes took place and diagnose problems.


See also:

- “Automatic Load Balancing for Failover Policy”
- “Load Balance Policy”
- “Path Verification”
- “PDO Removal”
- “Refreshing the Objects”

CLEARING PATH STATISTICS

You can Clear Path Statistics for all paths as needed for monitoring and diagnostic purposes.

To clear the statistics for ALL paths, do one of the following actions:

- From the Operations menu, choose Clear Path Statistics.
- In the Tree, right-click the LUN  icon, and choose Clear Path Statistics from the popup menu.

FEATURES AND SETTINGS

Features and Settings include the following functions:

- Automatic Load Balancing for Failover Policy
- Load Balance Policy
- Path Verification
- PDO Removal
- Performance Tab Refresh Rate
- Round Robin Count
- Refreshing the Objects
- Viewing System Information
- Saving System Information

AUTOMATIC LOAD BALANCING FOR FAILOVER POLICY

The PROMISE MPIO solution can load balance the paths for your LUNs with load balance policy set to Failover.

With Automatic Load Balancing enabled, the LUNs set to Failover policy are automatically redistributed among all available paths when:

- A path fails
- A failed path comes back online
- A new path is added

Automatic Load Balancing, when enabled, provides optimal data throughput for LUNs set to Failover policy.

Note that Automatic Load Balancing has NO effect upon LUNs set to Round Robin, Round Robin with Subset, or Least Queue Depth.

ENABLING AUTOMATIC LOAD BALANCING

To enable automatic load balancing:

From the Operations menu, choose Auto Load Balance.

When you see a check mark beside Auto Load Balance in the Operations menu, this feature is enabled.

See also:

- Viewing LUN Properties

- Viewing LUN Performance Statistics
- Load Balance Policy

LOAD BALANCE POLICY

Load Balance Policy is a method of equalizing the I/O traffic over each path by systematically dividing the load among multiple paths.

- Failover Policy – No load balancing. With Automatic Load Balancing disabled, the first path discovered is the primary path. I/Os follow the active path until it fails, then they change to next available path. Each LUN uses only one active path.

See Automatic Load Balancing for Failover Policy.

- Round Robin Policy – I/Os follow all active paths, changing paths at the specified I/O count. You can set the I/O count in the General tab of the Advanced Settings dialog box.

If LUN Affinity is enabled do NOT use Round Robin. See “Making Controller Settings” on page 59 or page 124.

- Round Robin with Subset Policy – One or more paths are designated as standby. I/Os follow all active paths, changing at the specified I/O count. You can set the I/O count in the General tab of the Advanced Settings dialog box.




If LUN Affinity is enabled, you can use Round Robin with Subset. See “Making Controller Settings” on page 59 or page 124.

- **Least Queue Depth Policy** – I/Os follow the path with the least number of requests queued.

Note that you can enable Automatic Load Balancing for LUNs with policy set to Failover. See Automatic Load Balancing for Failover Policy.

CHANGING LOAD BALANCE POLICY SETTINGS

To change load balance policy settings:

1. Do one of the following actions:
 - From the Operations menu, choose **Change Load Balance Policy**.
 - In Tree View, highlight a LUN  and click the Change Load Balance Policy  icon.
 - In Tree View, right-click a a LUN  and choose **Change Load Balance Policy** from the popup menu.

The Change Load Policy dialog box appears with the Load Balance Policy tab displayed.

2. Click the option button for one of the Load Policies.

- Failover Policy
- Round Robin Policy
- Round Robin with Subset Policy
- Least Queue Depth Policy

3. Click the Next button.

The Path Selection tab displays.

4. Take the action appropriate for your policy selection.

- For Round Robin and Least Queue Depth, no action is required. Skip to step 5.
- For Failover, move the path you want to be active to the **Primary Path Selected** pane. Move all other paths to the **Path Available** pane.
- For Round Robin with Subset, move the paths you want to be active to the **Primary Path Selected** pane. Move the paths you want as standby to the **Path Available** pane. You can have all paths in the Subset.

5. Click the Next button.

The Summary tab displays the current and selected (new) policy.

6. Click the Finish button to apply your settings.

The new settings take effect immediately.

See also:

- “Round Robin Count”
- “Refreshing the Objects”

PATH VERIFICATION

Path verification monitors any failed paths and automatically verifies them if they become available again.


There are two Path Verification Settings:

- Enable / disable
- Verification period in seconds

CHANGING PATH VERIFICATION SETTINGS

To make path verification settings:

1. Do one of the following actions:

- Click the Advanced Settings  icon.
- From the Operations menu, choose **Advanced Settings**.

The Advanced Settings dialog box appears with the MPIO Parameters tab displayed.

2. Check the Enable Path Verification box to enable path verification.

Uncheck to disable.

3. Click the arrows or type a new value in the Path Verification Period field to change the interval.

30 seconds is the default value.

4. Click the Apply button.

5. Click the OK button in the confirmation box.

The new setting takes effect immediately.

See also:

- Load Balance Policy
- “PDO Removal” on page 241
- “Refreshing the Objects” on page 242

PDO REMOVAL


PDO removal refers to the action of deleting a multipath input/output (MPIO) disk from the Windows Device Manager after all paths to a physical device object (PDO) have failed.

PDO removal interval refers to the period of time in seconds between the moment all paths to a PDO are disconnected and the MPIO disk disappears from the Device Manager.

CHANGING PDO REMOVAL SETTINGS

To change PDO removal settings:

1. Do one of the following actions:

- Click the Advanced Settings  icon.
- From the Operations menu, choose **Advanced Settings**.

The Advanced Settings dialog box appears with the MPIO Parameters tab displayed.

2. Click the arrows or type a new value in the PDO Remove Period field to change the interval.

120 seconds is the PROMISE-recommended default value.

- Click the Apply button.
- Click the OK button in the confirmation box.

The new setting takes effect immediately.


See also Path Verification.

PERFORMANCE TAB REFRESH RATE

Refresh Rate refers to the number of seconds between refreshes of the data reported on the Performance tab.

CHANGING REFRESH RATE SETTINGS

To change the refresh rate on the Performance tab:

- Do one of the following actions:
 - Click the Advanced Settings  icon.
 - From the Operations menu, choose **Advanced Settings**.
- The Advanced Settings dialog box appears with the MPIO Parameters tab displayed.
- Click the General tab.
- Under Refresh Rate, click the arrows or type a new value in the Seconds field to change the interval.

5 seconds is the default value.
- Click the Apply button.

The new setting takes effect immediately.

See also:



- Viewing LUN Performance Statistics
- Viewing Path Performance Statistics

ROUND ROBIN COUNT

When you set your path Load Balance Policy to Round Robin, the I/Os follow all active paths, changing paths at the specified I/O count. You can set the I/O count in the General tab of the Advanced Settings dialog box.

CHANGING THE ROUND ROBIN COUNT

To change Round Robin Count settings:

- Do one of the following actions:
 - Click the Advanced Settings  icon.
 - From the Operations menu, choose Settings  icon.
- The Advanced Settings dialog box appears with the MPIO Parameters tab displayed.
- Click the General tab.
- Under Round Robin Count, click the arrows or type a new value in the I/Os per Path field to change the count.

10 I/Os is the default value.
- Click the Apply button.

The new setting takes effect immediately.


See also:

- Load Balance Policy
- Viewing LUN Performance Statistics
- Viewing Path Performance Statistics

REFRESHING THE OBJECTS

Use this function after making an addition or deletion to your LUNs or paths.

To refresh the objects, do one of the following actions:

- From the Operations menu, choose Refresh.
- Click the Refresh  icon.


PerfectPath automatically displays all reported changes. However, some actions are not reported.

The Refresh action enables you to see the latest information.

- See also:
- Automatic Load Balancing for Failover Policy
 - Load Balance Policy
 - Path Verification
 - PDO Removal

VIEWING SYSTEM INFORMATION

To view System information and settings, do one of the following actions:

- From the System menu, choose System Information.
- Click the System Information  icon.

The System Information dialog box displays.

System information supplies information about the Host PC or Server, including:

Host Name	MPIO Version
Operating System	MPDEV File
OS Version	MPDEV Version
OS Manufacturer	MPSPFLTR File
IP Address	MPSPFLTR Version
Storport File	DSM File
Storport Version	DSM Version
MPIO File	




Note

File information includes the file name and location of the installed file in the server’s file system.

SAVING SYSTEM INFORMATION

To save the current System information and settings data to a text file:

1. Do one of the following actions:
 - From the System menu, choose System Information.
 - Click the System Information  icon.

The System Information dialog box displays.

2. From the System Information dialog box, click the Save button.
3. In the Save As dialog box, navigate to the folder where you want to save the file.
4. Type a file name into the File name field.

Append the file name with a .txt suffix.

5. Click the Save button.
6. Click the OK button in the confirmation box.

Your information and settings data are saved to a text file in the folder you designated.

TROUBLESHOOTING

If you see no LUNs in the PerfectPath GUI, or no Multi-Path Disk Devices under Disk drives (see Figure 3), do the following actions:

- Verify that there is at least one logical drive on the Vess R2600
- Check your HBA cards and driver installation
- Check your data connections

Make any needed corrections and reboot your Host PC as needed.

UPDATING PERFECTPATH

To update your PerfectPath software to the latest version:

1. Download the new PerfectPath installation file from PROMISE support:
<http://www.promise.com/support/> and save the installation file to your Windows desktop.
2. Manually remove the current PerfectPath installation.

See “Removing PerfectPath” on page 245.
3. Install the new PerfectPath software.

See Installing PerfectPath.

REPAIRING PERFECTPATH

To implement this procedure, you must use the same PerfectPath.exe installer file that you used to install the PerfectPath software onto your PC.

The installer’s version number is part of its file name. However, there is no corresponding number in the PerfectPath software.

To repair the PerfectPath software:

1. Double-click the PerfectPath.exe file to start the installer.
2. In the Welcome screen, click the Next button.
3. In the Program Maintenance screen, choose the Repair option, then click the Next button.
4. In the Ready to Repair the Program screen, click the Install button.
5. In the Install Completed screen, click the Finish button.
6. In the Restart message box, click the Yes button to restart your PC.

REMOVING PERFECTPATH

PREFERRED METHOD

This procedure uses the uninstaller included with PerfectPath.

To remove the PerfectPath software:

1. From the Start menu, choose All Programs > PerfectPath > Uninstall PerfectPath.
2. In the Welcome screen, click the Next button.
3. In the Program Maintenance screen, choose the Remove option, then click the Next button.
4. In the Remove the Program screen, click the Remove button.
5. In the Completed screen, click the Finish button.
6. In the Restart message box, click the Yes button to restart your PC.

ALTERNATE METHOD 1

This procedure uses the Windows uninstaller.

To remove the PerfectPath software:

1. In the Start menu, choose Control Panel, then choose Programs and Features.
2. Right-click Perfect Path, and choose Uninstall from the popup menu.
3. In the Confirmation box, click the Yes button.
4. In the Restart message box, click the Yes button to restart your PC.

ALTERNATE METHOD 2

To use this procedure, the PerfectPath.exe installer file must be the same version number as the PerfectPath software installed on your PC.

To remove the PerfectPath software:

1. Double-click the PerfectPath.exe file to start the installer.
2. In the Welcome screen, click the Next button.
3. In the Program Maintenance screen, choose the Remove option, then click the Next button.
4. In the Remove the Program screen, click the Remove button.

5. In the Completed screen, click the Finish button.
6. In the Restart message box, click the Yes button to restart your PC.

APPENDIX C: MULTIPATHING ON LINUX

The appendix covers the following topics:

- Before You Begin
- Task 1: Meeting Package Requirements
- Task 2: Preparing the Configuration File
- Task 3: Making Initial Host Settings
- Task 4: Create and Configure Devices
- Task 5: Setting-up ALUA
- RPM Packages and Documents for Linux MPIO
- Linux MPIO: Known Issues
- Sample multipath.conf File

PROMISE has fully tested Vess R2600 multipathing on RedHat RHEL 5.3, 5.4, and 5.5; and SuSE SLES 10 SP2, 10 SP3, 11, and 11 SP1. Coverage in this **Product Manual** is limited to those OSes.

Multipathing is possible on other Linux OSes. PROMISE has not tested every possible combination and therefore does not attempt to cover them here.

For a list of supported OSes, download the latest compatibility list from PROMISE support: <http://www.promise.com/support/>.

BEFORE YOU BEGIN

Before you can set up multipathing on your Linux Host PC, you must:

- Install your Fibre Channel or SAS HBA card into the Host PC.
- Install the HBA card drivers onto the Host PC
- Setup your Vess R2600, install your physical drives and create your logical drives.
- Attach your Fibre Channel or SAS cables from the HBA card to the Vess R2600 RAID subsystem.
- Install RHEL 5.x with the linux mpath option.
- For ALUA, refer to the PROMISE Linux support packages. See page 486.

Refer to the **Linux Administration Manual**, your HBA documentation, and this Appendix as needed for more information.

CHECK INITIAL SETUP

To check your initial setup, verify that you can view the logical drives on your Vess R2600 from your Linux desktop or terminal window. Refer to the **Linux Administration Manual** for the procedure on your system.

- If you can see your logical drives, the system is properly configured. Go to “Task 1: Meeting Package Requirements” on page 247.
- If you cannot see your logical drives, make the necessary adjustments and check again.

Task 1: Meeting Package Requirements

The latest device mapper and multipath packages must be loaded onto your Linux host before configuring Device Mapper Multipath (DM-MP). When this document was written, the current versions were:

- For RHEL 5.3
 - device-mapper-1.02.28-2.el5
 - device-mapper-multipath-0.4.7-23.el5
- For RHEL 5.4
 - device-mapper-1.02.32-1.el5
 - device-mapper-multipath-0.4.7-30.el5
- For RHEL 5.5
 - device-mapper-1.02.39-1.el5
 - device-mapper-multipath-0.4.7-34.el5
- For SLES 10 SP2
 - device-mapper-1.02.13-6.14
 - multipath-tools-0.4.7-34.38
- For SLES 10 SP3
 - device-mapper-1.02.13-6.14
 - multipath-tools-0.4.7-34.50.10
- For SLES 11
 - device-mapper-1.02.27-8.6
 - multipath-tools-0.4.8-40.1
- For SLES 11 SP1
 - device-mapper-1.02.27-8.20
 - multipath-tools-0.4.8-40.21.1

Installing Packages

The easiest and most effective way to install the device mapper and multipath tool is during OS installation. The device mapper installs by default, regardless of the configuration you select. However, you must manually specify the multipath tool, as it does not install as a part of any of the configurations of either OS. The multipath tool is listed as an option under Base System.

For hosts with the OS already installed, you can add the device mapper and multipath tool, if they are missing. Example: To add the multipath tool for RHEL 5.x, do the following actions:

1. Open a terminal window.
2. Type the following command and press **Enter**:

```
# rpm -ivh device-mapper-multipath-0.4.7-8.el5.i386.rpm
```

The system returns the following lines:

```
Preparing... ##### [100%]  
1:device-mapper-multipath##### [100%]  
#
```

(or a similar message)



Important

Where possible, obtain the device mapper and multipath tool from the original installation CDs to ensure full compatibility with your existing OS. Refer to your OS documentation for more information.

Verifying Packages – RedHat

To verify that the required packages are installed on the host, do the following actions:

1. Open a terminal window.
2. Type the following command and press **Enter**:

```
# rpm -qa | grep device-mapper
```

If the required packages are present, the system returns the following lines.

RHEL 5.3:

```
device-mapper-1.02.28-2.el5  
device-mapper-multipath-0.4.7-23.el5
```

RHEL 5.4:

```
device-mapper-1.02.32-1.el5  
device-mapper-multipath-0.4.7-30.el5
```

RHEL 5.5:

```
device-mapper-1.02.39-1.el5  
device-mapper-multipath-0.4.7-34.el5
```

Note that the actual version number might be different, depending on your configuration.

Verifying Packages – SuSE

To verify that the required packages are installed on the host, do the following actions:

1. Open a terminal window.
2. Type the following command and press **Enter**:

```
# rpm -qa | grep device-mapper
```

If the required package is present, the system returns the following line.

```
SLES 10 SP2:
device-mapper-1.02.13-6.14
```

```
SLES 10 SP3:
device-mapper-1.02.13-6.14
```

```
SLES 11:
device-mapper-1.02.27-8.6
```

```
SLES 11 SP1:
device-mapper-1.02.27-8.17.20
```

Note that the actual version number might be different, depending on your configuration.

3. Type the following command and press **Enter**:

```
# rpm -qa | grep multipath-tools
```

If the required package is present, the system returns the following line.

```
SLES 10 SP2:
multipath-tools-0.4.7-34.38
```

```
SLES 10 SP3:
multipath-tools-0.4.7-34.50.10
```

```
SLES 11:
multipath-tools-0.4.8-40.1
```

```
SLES 11 SP1:
multipath-tools-0.4.8-40.21.1
```

Note that the actual version number might be different, depending on your configuration.

Task 2: Preparing the Configuration File

To setup multipathing with Vess R2600, or any other subsystem, you must provide the required device attributes in a configuration file. The multipath configuration file is named multipath.conf. The functional version of the file is saved in the /etc directory.

RedHat Systems

For RedHat systems, there is a default /etc/multipath.conf file. However, the default file does not have the required device attributes to work with Vess R2600.

There are also sample configuration files in the /usr/share/doc/device-mapper-multipath-[version] directory:

- multipath.conf.annotated – multipath device attributes listed and defined
- multipath.conf.synthetic – multipath device attributes listed only

SuSE Systems

For SuSE systems, there is no default /etc/multipath.conf file.

There are sample multipath configuration files in the /usr/share/doc/packages/multipath-tools directory:

- multipath.conf.annotated – multipath device attributes listed and defined
- multipath.conf.synthetic – multipath device attributes listed only

Editing a Configuration File

You must provide a configuration file with required device attributes to work with Vess R2600.

Take the following actions to prepare a configuration file:

1. Choose an existing multipath.conf file and open the file in a text editor.
2. Save a working copy of the file under another name.
3. Edit the file to include the following line under defaults:

```
defaults {
    user_friendly_names yes
}
```

4. Edit the file to include the following lines under devices:

```
devices {
device {
    vendor      "Promise"
    product     "Vess R"
    path_grouping_policy    multibus
    getuid_callout    "/sbin/scsi_id -g -u -s /block/%n"
    path_checker      readsector0
    path_selector     "round-robin 0"
    hardware_handler  "0"
    failback          immediate
    rr_weight          uniform
    rr_min_io         100
    no_path_retry     20
    features          "1 queue_if_no_path"
    product_blacklist "Vess V-LUN"
}
}
```

5. Edit the file to include the following lines under devnode_blacklist:

```
devnode_blacklist {
devnode "^sda$"
devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
devnode "^hd[a-z]"
devnode "^cciss!c[0-9]d[0-9]*"
}
```

6. Be sure all the relevant lines of your configuration file are uncommented.

Remove the # character from the beginning of the line.

7. Save the file as multipath.conf.

8. Place a copy of the multipath.conf file into the Host’s /etc directory.

Task 3: Making Initial Host Settings

After the packages and configuration file are installed, the Host is ready to accept multipath settings.

Setting the Daemon to Run

RHEL 5.3, 5.4, 5.5

This action requires RHEL installation with the “linux mpath” Option. See Before You Begin.

To set the MP daemon to run at boot time:

1. Open a terminal window.
2. Verify that /etc/rc.d/rc[3~5].d/ has a symbolic link to /etc/rc.d/init.d/multipathd. Also see the Note below.

Run the command:

```
# ls -al /etc/rc.d/rc[3~5].d/ | grep multipathd
```

3. If there is no symbolic link, run the command:

```
# cd /etc/rc.d/rc[3~5].d/
```

Then run the command:

```
# ln -s /etc/rc.d/init.d/multipathd S13multipathd
```

Note: /etc/rc.d/rc[3~5].d saves a link to start at each run level.

- Run level 3 is for a single user.
- Run level 4 is for multiple users.
- Run level 5 is for multiple users on X Windows.

SLES 10 SP2, 10 SP3, 11, 11 SP1

To set the MP daemon to run at boot time:

1. Open a terminal window.
2. Set the daemon to run at boot time.

```
# chkconfig multipathd on
```

The system does not return anything.

VERIFYING THE MODULES ARE LOADED

To verify that the DM-MP modules are loaded:

1. Open a terminal window.
2. Verify that the multipath module is loaded.

```
# lsmod | grep dm_multipath
```

If the module is loaded, the system returns:

```
dm_multipath 215770 (or a similar message)
```

3. Verify that the device mapper module is loaded.

```
# lsmod | grep dm_mod
```

If the module is loaded, the system returns:

```
dm_mod 56537 8 dm_snapshot... (or a similar message)
```

VERIFYING THE DAEMON IS RUNNING

To verify that the MP daemon is running:

1. Open a terminal window.
2. Check the daemon’s status.

```
# /etc/init.d/multipathd status
```

3. Do one of the following actions:

- If the system returns:
multipathd is running (or a similar message)
Go to “Task 4: Create and Configure Devices” on page 480.
- If the system returns:
multipathd is stopped (or a similar message)
Start the MP daemon.

```
# /etc/init.d/multipathd start
```


Then go to Task 4: Create and Configure Devices.

TASK 4: CREATE AND CONFIGURE DEVICES

This step applies the settings from the multipath.config file to the Host.

1. Open a terminal window.
2. Type the following command and press **Enter**:

```
#multipath -v3
```

The system returns:

```
...
==== paths list ====

uuid          hcil   dev dev_t pri dm_st chk_st  vend...
222490001555459b3b 2:0:0:0 sdb 8:16  1  [undef][undef] Prom...
222b40000155a75b49 2:0:0:2 sbc 8:32  1  [undef][undef] Prom...
20efcff5501000121a 3:0:0:0 sbd 8:48  1  [undef][undef] Prom...
222b44000155ebf0c  3:0:0:1 sde 8:64  1  [undef][undef] Prom...
params = 1 que_if_no_path 0 1 1 round-robin 0 1 1 8:64 100
status = 1 0 0 1 1 A 0 1 0 8:64 A 0
sde: mask = 0x4
sde: path checker = readsector0 (controller setting)
sde: state = 2
...
```

(or a similar message)

3. Restart the MP daemon.

```
#/etc/init.d/multipathd restart
```

The system returns:

```
Stopping multipathd daemon (or a similar message)
```

```
Starting multipathd daemon (or a similar message)
```

For more information about path monitoring functions, type help and press **Enter**.

Task 5: Setting-up ALUA

Vess R2600 supports Asymmetric Logical Unit Access (ALUA) on the latest Linux distributions:

- RedHat Linux RHEL 5.4
- RedHat Linux RHEL 5.5
- SuSE Linux SLES 10 SP3
- SuSE Linux SLES 11
- SuSE Linux SLES 11 SP1

PROMISE provides RPM packages and multipath.conf files for ALUA. See “RPM Packages and Documents for Linux MPIO” on page 486.

RedHat Linux RHEL 5.4

KERNEL WITH XEN OR PAE SUPPORT

To support ALUA within RHEL 5.4 using kernel with XEN or PAE support:

1. Install RHEL 5.4. When the CD is chosen for installation, immediately type linux mpath.
2. Copy the PROMISE-supplied multipath.conf file to the /etc directory.
cp multipath.conf-Vess R2600-alua-rhel5-4 /etc/multipath.conf

3. Install this patched alua_dh_scsi source module.
rpm -ivh scsi_dh_alua_Vess R2600Dual-1-1.src.rpm

4. Make your own rpm.
cd /usr/src/redhat/SPEC
rpmbuild -ba scsi_dh_alau.spec

5. Determine your system architecture, x86_64 or i386.
uname -a

6. Install the appropriate scsi_dh_alua driver module.
 - i386
cd /usr/src/redhat/RPM/i386
rpm -ivh scsi_dh_alua_Vess R2600-1-1.i386.rpm
 - x86_64
cd /usr/src/redhat/RPM/x86_64

rpm -ivh scsi_dh_alua_Vess R2600-1-1.x86_64.rpm

7. Check the multipath configuration.
Multipath -ll (multipath -v4)

RedHat Linux RHEL 5.5

KERNEL WITH XEN OR PAE SUPPORT

To support ALUA within RHEL 5.5 using kernel with XEN or PAE support:

1. Install RHEL 5.5. When the CD is chosen for installation, immediately type linux mpath.
2. Copy the PROMISE-supplied multipath.conf file to the /etc directory.
cp multipath.conf.alua-rhel5.5 /etc/multipath.conf

3. Install this patched alua_dh_scsi driver module.
4. rpm -ivh scsi_dh_alua_Vess R2600Dual-2-1.src.rpm

5. Make your own rpm.
cd /usr/src/redhat/SPEC
rpmbuild -ba scsi_dh_alau.spec

6. Determine your system architecture, x86_64 or i386.
uname -a

7. Install the appropriate scsi_dh_alua driver module.
 - i386
cd /usr/src/redhat/RPM/i386
rpm -ivh scsi_dh_alua_Vess R2600-2-1.i386.rpm
 - x86_64
cd /usr/src/redhat/RPM/x86_64
rpm -ivh scsi_dh_alua_Vess R2600-2-1.x86_64.rpm

8. Check the multipath configuration.
Multipath -ll (multipath -v4)

SuSE Linux SLES 10 SP3

To support ALUA within SLES10 SP3:

1. Determine your system architecture, i586 or x86_64.

```
#> uname -a
```

2. Do one of the following actions:

- Remove the currently installed multipath tool.

```
rpm -ev multipath-tools
```

Install the appropriate new multipath-tool RPM package.

i586

```
rpm -ivh multipath-tools-0.4.7-34.50.10.ass.fix.i586.rpm
```

x86_64

```
rpm -ivh multipath-tools-0.4.7-34.50.10.ass.fix.x86_64.rpm
```

- Force install the appropriate new multipath-tool RPM package.

i586

```
#> rpm -ivh -force multipath-tools-0.4.7-34.50.10.ass.fix.i586.rpm
```

x86_64

```
#> rpm -ivh -force multipath-tools-0.4.7-34.50.10.ass.fix.x86_64.rpm
```

3. Copy the PROMISE-supplied multipath.conf file to the /etc directory.

```
#> cp multipath.conf-Vess R2600-alua-sles10-sp3 /etc multipath.conf
```

4. Check the multipath configuration.

```
chkconfig multipathd on
```

chkconfig multipathd (shows the status of multipathd)

```
chkconfig boot.multipath on
```

chkconfig boot.multipath (shows the status of boot.multipath)

5. Reboot the Vess R2600.

6. Check the multipath configuration.

```
Multipath -ll (multipath -v4)
```

SUSE LINUX SLES 11

To support ALUA within SLES 11:

1. Determine your system architecture, i586 or x86_64.

```
#> uname -a
```

2. Install the appropriate path priority tool library.

- i586

```
#> rpm -ivh multipath-promise-suse11-0.4.8-1.i586.rpm
```

- x86_64

```
#> rpm -ivh multipath-promise-suse11-0.4.8-1.x86_64.rpm
```

3. Copy the PROMISE-supplied multipath.conf file to the /etc directory.

```
#> cp multipath.conf-Vess R2600-alua-sles11 /etc multipath.conf
```

4. Check the multipath configuration.

```
chkconfig multipathd on
```

chkconfig multipathd (shows the status of multipathd)

5. Reboot the Vess R2600.

6. Check the multipath configuration.

```
Multipath -ll (multipath -v4)
```

SUSE LINUX SLES 11 SP1

Not updated with Novell SP1 patches

If you did NOT update with Novell SLES 11 SP1 patches:

1. Determine your system architecture, i586 or x86_64.

```
#> uname -a
```

2. Install the appropriate patched scsi_dh_alua.ko device handler.

- i586

```
#> rpm -ivh --force scsi_dh_alua_sles11sp1-2-1.i586.rpm
```

- x86_64

```
#> rpm -ivh --force scsi_dh_alua_sles11sp1-2-1.x86_64.rpm
```

3. Copy the PROMISE-supplied multipath.conf file to the /etc directory.

```
#> cp multipath.conf.alua-sles11sp1 /etc multipath.conf
```

4. Check the multipath configuration.

```
chkconfig multipathd on
```

chkconfig multipathd (shows the status of multipathd)

5. Reboot the Vess R2600.

6. Check the multipath configuration.

```
Multipath -ll (multipath -v4)
```

UPDATED WITH NOVELL SP1 PATCHES

If you updated with Novell SLES 11 SP1 patches:

1. Copy the PROMISE-supplied multipath.conf file to the /etc directory.
#> cp multipath.conf.alua-sles11sp1 /etc multipath.conf
2. Check the multipath configuration.
chkconfig multipathd on
chkconfig multipathd (shows the status of multipathd)
3. Reboot the Vess R2600.
4. Check the multipath configuration.
Multipath -ll (multipath -v4)

FIBRE CHANNEL HBA TO VESS R2600

If you plan to use connect a Fibre Channel HBA card to Vess R2600 RAID subsystem, you must change the remote port (rport) configuration.

1. Change the dev_loss_tmo value as large as possible.
Example: 0x7fffffff.
2. Change the fast_io_fail_tmo value as 30.

RPM PACKAGES AND DOCUMENTS FOR LINUX MPIO

PROMISE provides RPM packages and multipath.conf files for Linux on the support: <http://www.promise.com/support/>.

A PROMISE Linux package contains:

- package folder – RPM packages for the Linux OS
- multipath-conf folder
 - ALUA folder – Configuration file and instructions for ALUA, if supported
 - normal folder – Configuration file for general multipathing



Important

Please read the **How to Configure** document in the ALUA folder for the latest information before beginning your setup.

The table below lists the content of each PROMISE Linux package:

Package	Folder	Contents
SLES10-SP2	normal	multipath.conf-Vess R2600-normal-sles10-sp2
SLES10-SP3	package	multipath-tools-0.4.7-34.50.10.ass.fix.i586.rpm multipath-tools-0.4.7-34.50.10.ass.fix.x86_64.rpm
	ALUA	multipath-conf-Vess R2600-alua-sles10-sp3 How to Configure MPIO SLES10 SP3.doc
	normal	multipath-conf-Vess R2600-normal-sles10-sp3
SLES11	packages	multipath-promise-suse11-0.4.8-i586.rpm multipath-promise-suse11-0.4.8-x86_64.rpm
	ALUA	multipath-conf-Vess R2600-alua-sles11 How to Configure MPIO SLES 11.doc
	normal	multipath-conf-Vess R2600-normal-sles11
SLES11-SP1	packages	scsi_dh_alua_sles11sp1-2-1.i586.rpm scsi_dh_alua_sles11sp1-2-1.x86_64.rpm scsi_dh_alua_sles11sp1-2-1.src.rpm
	ALUA	multipath.conf.alua-sles11sp1 How to Configure MPIO SLES 11sp1.doc
	normal	multipath.conf-Vess R2600-normal-sles11sp1 How to Configure MPIO SLES 11sp1.doc
RHEL-5.3	normal	multipath.conf-Vess R2600-normal-rhel5-3

RHEL-5.4	package	scsi_dh_alua_Vess R2600Dual -1-1.src.rpm
	ALUA	multipath.conf-Vess R2600-alua-rhel5-4 How to Configure MPIO RHEL 5.4.doc
	normal	multipath.conf-Vess R2600-normal-rhel5-4
RHEL-5.5	packages	scsi_dh_alua_Vess R2600Dual- 2-1.src.rpm
	ALUA	multipath.conf.alua-rhel5.5 How to Configure MPIO RHEL 5.5.doc
	normal	multipath.conf-normal-rhel5.5

LINUX MPIO: KNOWN ISSUES

Issue	OS	Description
1	SLES 11 GMC	OS multipath with SAS interface makes kernel panic.
2	SLES 11 GMC	OS sets default rports dev_loss_tmo value at 10 seconds resulting in loss of the path during failover/failback.
3	RHEL 5.4	With FC switch, HBA driver (such as Emulex 4g, QLogic 8g) sets rports dev_loss_tmo value too small, resulting in loss of the path during failover/failback. To fix this problem, increase the value to 60 seconds.
4	SLES 10 SP3	With FC switch, HBA driver (such as Emulex 4g, QLogic 8g) sets rports dev_loss_tmo value too small, resulting in loss of the path during failover/failback. To fix this problem, increase the value to 60 seconds.
5	SLES 11 SP1	With a FC HBA, set rports dev_loss_tmo as large as possible, such as 0x7ffffff. Set fast_io_failure_tmo to 30 seconds. Without these settings, the system does not recognize device or the systems hang during boot and failover/failback. Download and update the latest SLES11 SP1 patches from Novell. Without these patches, the system hangs during failover/failback.

With a Fibre Channel switch, a dev_loss_tmo value set too small can result in loss of the path during failover/failback.

Linux maintains the rports dev_loss_tmo value in the /sys/class/fc_remote_port/rport-xxxx/dev_loss_tmo file. The Fibre Channel HBA driver sets this value at loading time.

To change the dev_loss_tmo value to the recommended 60 seconds, during runtime type the echo command:
echo 60 > /sys/class/fc_remote_port/rpot-1:0:0/dev_loss_tmo

SAMPLE MULTIPATH.CONF FILE

Below is a complete multipath.conf file for Vess R2600.

- If you have no other multipath devices on your Host, you can use this multipath.conf file as shown.
- If you have other multipath devices, add these settings to your existing multipath.conf file.

```
##
## This is a template multipath-tools configuration file
## for the Promise Vess R2600 subsystem
##
defaults {
    user_friendly_names yes
}
blacklist {
    devnode "^sda$"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z][[0-9]*]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices {
    device {
        vendor          "Promise"
        product          "Vess R2600"
        path_grouping_policy multibus
        getuid_callout    "/sbin/scsi_id -g -u -s /block/%n"
        path_checker      readsector0
        path_selector      "round-robin 0"
        hardware_handler  "0"
        failback          immediate
        rr_weight          uniform
        rr_min_io          100
        no_path_retry      20
        features           "1 queue_if_no_path"
        product_blacklist "Vess R2600 V-LUN"
    }
}
```



2F, No. 30, Industry E. Rd. IX, Science-Based Industrial Park, Hsinchu 30075, Taiwan, R.O.C.
Tel: +886-3-5782-395 Fax: +886-3-5782-390 www.promise.com